



SYCLTM Specification

Generic heterogeneous computing for modern C++

Version 2020 provisional

Document Revision: 1

Revision Date: June 30, 2020

Git revision: tags/SYCL-2020/provisional-rev1-0-gfaaae8e88

Khronos[®] SYCLTM Working Group

Editors: Ronan Keryell, Maria Rovatsou & Lee Howes

Copyright 2011-2020 The Khronos® Group, Inc. All Rights Reserved

Copyright© 2011-2020 The Khronos® Group, Inc. All Rights Reserved.

This specification is protected by copyright laws and contains material proprietary to the Khronos® Group, Inc. It or any components may not be reproduced, republished, distributed, transmitted, displayed, broadcast, or otherwise exploited in any manner without the express prior written permission of Khronos® Group. You may use this specification for implementing the functionality therein, without altering or removing any trademark, copyright or other notice from the specification, but the receipt or possession of this specification does not convey any rights to reproduce, disclose, or distribute its contents, or to manufacture, use, or sell anything that it may describe, in whole or in part.

Khronos[®] Group grants express permission to any current Promoter, Contributor or Adopter member of Khronos[®] to copy and redistribute UNMODIFIED versions of this specification in any fashion, provided that NO CHARGE is made for the specification and the latest available update of the specification for any version of the API is used whenever possible. Such distributed specification may be reformatted AS LONG AS the contents of the specification are not changed in any way. The specification may be incorporated into a product that is sold as long as such product includes significant independent work developed by the seller. A link to the current version of this specification on the Khronos[®] Group website should be included whenever possible with specification distributions.

Khronos[®] Group makes no, and expressly disclaims any, representations or warranties, express or implied, regarding this specification, including, without limitation, any implied warranties of merchantability or fitness for a particular purpose or non-infringement of any intellectual property. Khronos[®] Group makes no, and expressly disclaims any, warranties, express or implied, regarding the correctness, accuracy, completeness, timeliness, and reliability of the specification. Under no circumstances will the Khronos[®] Group, or any of its Promoters, Contributors or Members or their respective partners, officers, directors, employees, agents, or representatives be liable for any damages, whether direct, indirect, special or consequential damages for lost revenues, lost profits, or otherwise, arising from or in connection with these materials.

Khronos® is a registered trademark and SYCLTM, SPIRTM, WebGLTM, EGLTM, COLLADATM, StreamInputTM, OpenVXTM, OpenKCamTM, glTFTM, OpenKODETM, OpenVGTM, OpenWFTM, OpenSL ESTM, OpenMAXTM, OpenMAX ALTM, OpenMAX ILTM and OpenMAX DLTM and WebCLTM are trademarks of the Khronos® Group Inc. OpenCLTM is a trademark of Apple Inc. and OpenGL® and OpenML® are registered trademarks and the OpenGL ESTM and OpenGL SCTM logos are trademarks of Silicon Graphics International used under license by Khronos®. All other product names, trademarks, and/or company names are used solely for identification and belong to their respective owners.

Contents

| 1 | Ackı | owledgements | 25 |
|---|-------|--|----|
| 2 | Intro | duction | 27 |
| 3 | SYC | architecture | 29 |
| | 3.1 | Overview | 29 |
| | 3.2 | Anatomy of a SYCL application | 30 |
| | 3.3 | Normative references | 31 |
| | 3.4 | The SYCL platform model | 31 |
| | 3.5 | The SYCL backend model | 32 |
| | | 3.5.1 Platform mixed version support | 33 |
| | 3.6 | SYCL execution model | 33 |
| | | 3.6.1 SYCL application execution model | 34 |
| | | 3.6.1.1 SYCL backend resources managed by the SYCL application | 34 |
| | | 3.6.1.2 SYCL command groups and execution order | 35 |
| | | 3.6.2 SYCL kernel execution model | 36 |
| | 3.7 | Memory model | 37 |
| | 011 | 3.7.1 SYCL application memory model | 37 |
| | | 3.7.2 SYCL device memory model | 40 |
| | | 3721 Access to memory | 41 |
| | | 3722 Memory consistency inside SYCL kernels | 41 |
| | | 3723 Atomic operations | 43 |
| | 38 | The SVCL programming model | 43 |
| | 5.0 | 3.8.1 Minimum version of C++ | 43 |
| | | 3.8.2 Alignment with future versions of C++ | 43 |
| | | 3.8.3 Basic data parallel kernels | 43 |
| | | 3.8.4 Work group data parallal karnels | |
| | | 3.8.5 Hierarchical data parallel kernels | 44 |
| | | 3.6. Vernels that are not lounched over perallel instances | 44 |
| | | 3.8.7 Pre defined kernels | 45 |
| | | 2.9.9 Symphronization | 45 |
| | | 2.9.9.1 Symphysization in the SVCL application | 43 |
| | | 3.8.8.1 Synchronization in SVCL learnels | 43 |
| | | 5.0.0.2 Synchronization in STCL kennels | 40 |
| | | 5.8.9 Error nandling | 40 |
| | | 5.8.10 Failback mechanism | 40 |
| | | 5.8.11 Scheduling of kernels and data movement | 47 |
| | | 3.8.12 Managing object lifetimes | 47 |
| | | 3.8.13 Device discovery and selection | 48 |
| | 2.0 | 3.8.14 Interfacing with SYCL backend API | 48 |
| | 3.9 | Memory objects | 48 |
| | 3.10 | SYCL device compiler | 49 |
| | | 3.10.1 Building a SYCL program | 49 |
| | | 3.10.2 Naming of kernels | 50 |

| | 3.11 | Langu | age restrict | ions in kernels | 50 |) |
|---|------|---------|--------------|---|-----|----------|
| | | 3.11.1 | SYCL lir | ker | 51 | L |
| | | 3.11.2 | Function | and data types available in kernels | 51 | L |
| | 3.12 | Endiar | ness supp | vrt | 51 | L |
| | 3.13 | Examp | ole SYCL a | pplication | 51 | L |
| | ~~~~ | - | | | | |
| 4 | SYC | CL prog | ramming | nterface | 55 | ; |
| | 4.1 | Backer | nds | | 55 | <u>,</u> |
| | | 4.1.1 | Backend | macros | 55 | <u>,</u> |
| | 4.2 | Generi | c vs non-g | eneric SYCL | 55 | 5 |
| | 4.3 | Header | r files and | namespaces | 56 | 5 |
| | 4.4 | Class a | availability | | 56 | 5 |
| | 4.5 | Comm | on interfac | e | 57 | ! |
| | | 4.5.1 | Param tra | its class | 57 | ! |
| | | 4.5.2 | Backend | interoperability | 57 | / |
| | | | 4.5.2.1 | Type traits backend_traits | 57 | / |
| | | | 4.5.2.2 | Template function get_native | 58 | 5 |
| | | | 4.5.2.3 | Template functions make_* | 58 | 3 |
| | | 4.5.3 | Common | reference semantics | 60 |) |
| | | 4.5.4 | Common | by-value semantics | 62 | 2 |
| | | 4.5.5 | Propertie | 3 | 64 | Ł |
| | | | 4.5.5.1 | Properties interface | 64 | Ł |
| | 4.6 | SYCL | runtime cl | asses | 66 | 5 |
| | | 4.6.1 | Device se | lection | 66 | 5 |
| | | | 4.6.1.1 | Device selector | 66 | 5 |
| | | 4.6.2 | Platform | class | 68 | 3 |
| | | | 4.6.2.1 | Platform interface | 69 |) |
| | | | 4.6.2.2 | Platform information descriptors | | L |
| | | 4.6.3 | Context of | lass | 72 | 2 |
| | | | 4.6.3.1 | Context interface | 72 | 2 |
| | | | 4.6.3.2 | Context information descriptors | 74 | Ł |
| | | | 4.6.3.3 | Context properties | 75 | 5 |
| | | 4.6.4 | Device cl | ass | 75 | 5 |
| | | | 4.6.4.1 | Device interface | | 5 |
| | | | 4.6.4.2 | Device information descriptors | |) |
| | | | 4.6.4.3 | Device aspects | 91 | L |
| | | 4.6.5 | Queue cl | 188 | 94 | Ł |
| | | | 4.6.5.1 | Queue interface | 95 | 5 |
| | | | 4.6.5.2 | Queue information descriptors | 103 | 3 |
| | | | 4.6.5.3 | Queue properties | 103 | 3 |
| | | | 4.6.5.4 | Queue error handling | 104 | Ł |
| | | 4.6.6 | Event cla | 58 | 105 | 5 |
| | | | 4.6.6.1 | Event information and profiling descriptors | 108 | 3 |
| | 4.7 | Data a | ccess and s | torage in SYCL | 109 |) |
| | | 4.7.1 | Host allo | cation | 109 |) |
| | | | 4.7.1.1 | Default allocators | 109 |) |
| | | 4.7.2 | Buffers | | 110 |) |
| | | | 4.7.2.1 | Buffer interface | 111 | Ĺ |
| | | | 4.7.2.2 | Buffer properties | 121 | Ĺ |
| | | | 4.7.2.3 | Buffer synchronization rules | 122 | , |
| | | 4.7.3 | Images | · · · · · · · · · · · · · · · · · · · | 124 | Ł |
| | | | 0 | | | |

CONTENTS

| | | 4.7.3.1 Unsampled image interface |
|-----|---------|---|
| | | 4.7.3.2 Sampled image interface |
| | | 4.7.3.3 Image properties |
| | | 4.7.3.4 Image synchronization rules |
| | 4.7.4 | Sharing host memory with the SYCL data management classes |
| | | 4.7.4.1 Default behavior |
| | | 4.7.4.2 SYCL ownership of the host memory |
| | | 4.7.4.3 Shared SYCL ownership of the host memory |
| | 4.7.5 | Synchronization primitives |
| | 4.7.6 | Accessors |
| | | 4.7.6.1 Access targets |
| | | 4.7.6.2 Access modes |
| | | 4.7.6.3 Access tags |
| | | 4.7.6.4 Device and host accessors |
| | | 4.7.6.5 Placeholder accessor |
| | | 4.7.6.6 Accessor declaration |
| | | 4.7.6.7 Constness of the accessor data type |
| | | 4.7.6.8 Implicit accessor conversions |
| | | 4.7.6.9 Device buffer accessor |
| | | 4.7.6.9.1 Device buffer accessor interface |
| | | 4.7.6.9.2 Device buffer accessor properties |
| | | 4.7.6.10 Host buffer accessor |
| | | 4.7.6.10.1 Host buffer accessor interface |
| | | 4.7.6.10.2 Host buffer accessor properties |
| | | 4.7.6.11 Local accessor |
| | | 4.7.6.11.1 Local accessor interface |
| | | 4.7.6.11.2 Local accessor properties |
| | | 4.7.6.12 Image accessor |
| | | 4.7.6.12.1 Image accessor interface |
| | | 4.7.6.12.2 Image accessor properties |
| | 4.7.7 | Address space classes |
| | | 4.7.7.1 Multi-pointer class |
| | | 4.7.7.2 Explicit pointer aliases |
| | 4.7.8 | Samplers |
| 4.8 | Unified | l shared memory (USM) |
| | 4.8.1 | Unified addressing |
| | 4.8.2 | Kinds of unified shared memory |
| | | 4.8.2.1 Explicit USM |
| | | 4.8.2.2 Restricted USM |
| | | 4.8.2.3 Concurrent USM |
| | | 4.8.2.4 System USM |
| | 4.8.3 | USM allocations |
| | 4.8.4 | C++ allocator interface |
| | 4.8.5 | Utility functions |
| | | 4.8.5.1 Explicit USM |
| | | 4.8.5.1.1 malloc_device |
| | | 4.8.5.1.2 aligned_alloc_device |
| | | 4.8.5.1.3 memcpy |
| | | 4.8.5.1.4 memset |
| | | 4.8.5.1.5 fill |
| | | 4.8.5.2 Restricted USM |

| | | 4.8.5.2.1 malloc |
|------|------------|--|
| | | 4.8.5.2.2 aligned_alloc_host |
| | | 4.8.5.2.3 Performance hints |
| | | 4.8.5.2.3.1 prefetch |
| | | 4.8.5.3 Concurrent USM |
| | | 4.8.5.3.1 Performance hints |
| | | 4.8.5.3.1.1 prefetch |
| | | 4.8.5.3.1.2 mem_advise |
| | | 4.8.5.4 General |
| | | 4.8.5.4.1 malloc |
| | | 4.8.5.4.2 aligned alloc |
| | | 4.8.5.4.3 free |
| | 486 | Unified shared memory information 211 |
| | 1.0.0 | 4861 Pointer queries 211 |
| | | $48611 \text{aet pointer type} \qquad 211$ |
| | | $\frac{48612}{12} \text{ get pointer device} $ |
| 4.0 | SVCI | $4.6.0.1.2 \text{get_pointer_device} \dots \dots \dots \dots \dots \dots \dots \dots \dots $ |
| 4.9 | A 0 1 | DACs without second and a second |
| | 4.9.1 | |
| 4.10 | 4.9.2 E | |
| 4.10 | Expres | sing parallelism through kernels |
| | 4.10.1 | Ranges and index space identifiers |
| | | 4.10.1.1 range class |
| | | 4.10.1.2 nd_range class |
| | | 4.10.1.3 id class |
| | | 4.10.1.4 item class |
| | | 4.10.1.5 nd_item class |
| | | 4.10.1.6 h_item class |
| | | 4.10.1.7 group class |
| | | 4.10.1.8 sub_group class |
| | 4.10.2 | Reduction variables |
| | | 4.10.2.1 reduction interface |
| | | 4.10.2.2 reducer class |
| | 4.10.3 | Command group scope |
| | 4.10.4 | Command group handler class |
| | 4.10.5 | Class kernel_handler |
| | | 4.10.5.1 Constructors |
| | | 4.10.5.2 Member functions |
| | 4.10.6 | SYCL functions for adding requirements |
| | 4.10.7 | SYCL functions for invoking kernels |
| | | 4.10.7.1 single task invoke |
| | | 4 10 7 2 parallel for invoke 253 |
| | | 4 10 7 3 Parallel for hierarchical invoke |
| | 4 10 8 | SYCI functions for explicit memory operations 260 |
| | 4.10.0 | Functions for using a module 262 |
| | 4 10 10 | Functions for using specialization constants 262 |
| 1 11 | Host to | she 262 |
| 4.11 | 110St ta | Overview 242 |
| | 4.11.1 | Class interes handle 203 |
| | 4.11.2 | Class Interop_nancie 204 4.11.2.1 Constructors 205 |
| | | 4.11.2.1 CONSTRUCTORS |
| | 4.1.1.2 | 4.11.2.2 Template member functions get_native_* |
| | 4.11.3 | Additions to the handler class |

CONTENTS

| | Kernel | class | | 267 |
|--|--|--|---|---|
| | 4.12.1 | Kernel inf | ormation descriptors | 269 |
| 4.13 | Modul | es | | . 271 |
| | 4.13.1 | Overview | | . 271 |
| | 4.13.2 | Specializa | tion constants | . 272 |
| | 4.13.3 | Synopsis | | 272 |
| | 4.13.4 | Enum clas | smodule_state | . 274 |
| | 4.13.5 | Class tem | late specialization_id | 275 |
| | | 4.13.5.1 | Constructors | 275 |
| | | 4.13.5.2 | Special member functions | 275 |
| | 4.13.6 | Class tem | late module | 275 |
| | | 4.13.6.1 | Constructors | 276 |
| | | 41362 | Member functions | 276 |
| | 4 13 7 | Free funct | | 278 |
| | 4 13 8 | Namesnac | ethis module | 280 |
| | 4.15.0 | 4 13 8 1 | Type traits | 280 |
| | | 4.13.0.1 | Free functions | 280 |
| 1 1 1 | Dofinir | 4.13.0.2 | | 200 |
| 4.14 | | ig kernels. | · · · · · · · · · · · · · · · · · · · | 201 |
| | 4.14.1 | | | . 281 |
| | 4.14.2 | Defining I | ernels as lambda functions | . 282 |
| | 4.14.3 | Defining I | ernels using modules | . 283 |
| | 4.14.4 | Rules for | parameter passing to kernels | . 284 |
| 4.15 | Error h | andling | | . 285 |
| | 4.15.1 | Error hand | ling rules | 285 |
| | | 4.15.1.1 | Asynchronous error handler | . 285 |
| | | 4.15.1.2 | Behavior without an async_handler | . 285 |
| | | 4.15.1.3 | Priorities of async_handlers | . 285 |
| | | | A superior on a superior with a secondary quare | |
| | | 4.15.1.4 | | 286 |
| | 4.15.2 | 4.15.1.4 Exception | | 286 287 |
| 4.16 | 4.15.2 Data ty | 4.15.1.4 Exception | class interface | 286 287 292 |
| 4.16 | 4.15.2 Data ty 4.16.1 | 4.15.1.4 Exception pes Scalar dat | | 286 287 292 292 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception pes Scalar dat Vector typ | class interface | 286 287 292 292 292 292 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 | Asynchronous errors with a secondary queue class interface | 286 287 292 292 292 292 293 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 | Asynchronous enors with a secondary queue class interface types es Vec interface Aliases | 286 287 292 292 292 292 293 310 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 | Asynchronous enors with a secondary queue class interface a types es vec interface Aliases Swizzles | 286 287 292 292 292 293 310 310 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 | Asynchronous errors with a secondary queue class interface types es vec interface Aliases Swizzles Swizzled vec class | 286 287 292 292 292 293 310 310 311 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 4.16.2.5 | Asynchronous enfors with a secondary queue class interface a types es vec interface Aliases Swizzles Swizzled vec class Rounding modes | 286 287 292 292 293 310 311 311 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.5 | Asynchronous errors with a secondary queue class interface a types es vec interface Aliases Swizzles Swizzled vec class Rounding modes Memory layout and alignment | 286 287 292 292 292 293 310 310 311 311 311 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.6 4.16.2.7 | Asynchronous enors with a secondary queue class interface a types es vec interface Aliases Swizzles Swizzled vec class Rounding modes Memory layout and alignment Considerations for endianness | 286 287 292 292 292 293 310 311 311 311 312 312 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.5 4.16.2.6 4.16.2.7 4.16.2.8 | Asynchronous enors with a secondary queue class interface a types es vec interface Aliases Swizzles Swizzled vec class Rounding modes Memory layout and alignment Considerations for endianness Performance note | 286 287 292 292 292 293 310 310 311 311 312 312 312 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.6 4.16.2.7 4.16.2.8 Math arra | Asynchronous enors with a secondary queue class interface a types es wec interface Aliases Swizzles Swizzled vec class Rounding modes Memory layout and alignment Considerations for endianness Performance note | 286 287 292 292 292 293 310 310 311 311 312 312 312 312 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.6 4.16.2.7 4.16.2.8 Math arra 4.16.3.1 | Asynchronous enors with a secondary queue class interface a types es vec interface Aliases Swizzles Swizzles Swizzled vec class Rounding modes Memory layout and alignment Considerations for endianness Performance note ytypes Math array interface | 286 287 292 292 293 310 311 311 312 312 312 312 312 312 312 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.6 4.16.2.7 4.16.2.8 Math arra 4.16.3.1 4.16.3.2 | Asynchronous enors with a secondary queue class interface a types es we interface Aliases Swizzles Swizzled vec class Rounding modes Memory layout and alignment Considerations for endianness Performance note 'types Math array interface | 2866 2877 2922 2922 2933 3100 3111 3111 312 312 312 312 312 312 312 3 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.6 4.16.2.7 4.16.2.8 Math arra 4.16.3.1 4.16.3.2 4.16.3.3 | Asynchronous enors with a secondary queue class interface | 2866 2877 2922 2922 2933 310 311 311 312 312 312 312 312 312 312 312 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 4.16.3 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.6 4.16.2.7 4.16.2.8 Math arra 4.16.3.1 4.16.3.2 4.16.3.3 | Asynchronous enors with a secondary queue class interface a types es vec interface Aliases Swizzles Swizzles Swizzles Considerations for endianness Performance note 'types Math array interface Aliases Memory layout and alignment of types Math array interface Aliases Memory layout and alignment of types Math array interface Aliases Memory layout and alignment of atomics | 2866 2877 2922 2922 2933 3100 3111 312 312 312 312 312 312 312 312 323 323 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 4.16.3 Synchr 4.17 1 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.6 4.16.2.7 4.16.2.8 Math arra 4.16.3.1 4.16.3.2 4.16.3.3 conization a | Asynchronous enors with a secondary queue class interface a types es vec interface Aliases Swizzles Swizzled vec class Rounding modes Memory layout and alignment Considerations for endianness Performance note 'types Math array interface Aliases Memory layout and alignment of the forces | 28662 2877 2922 2922 2933 310 311 311 312 312 312 312 312 312 312 323 323 |
| 4.164.17 | 4.15.2 Data ty 4.16.1 4.16.2 4.16.3 Synchr 4.17.1 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.6 4.16.2.7 4.16.2.8 Math arra 4.16.3.1 4.16.3.2 4.16.3.3 conization a Barriers a | Asynchronous errors with a secondary queue class interface class interface es vec interface Aliases Swizzles Swizzled vec class Rounding modes Memory layout and alignment Considerations for endianness Performance note 'types Math array interface Aliases Memory layout and alignment of the face in the face | 2866 2877 2922 2922 2933 3100 3111 312 312 312 312 312 312 312 312 323 323 |
| 4.164.17 | 4.15.2 Data ty 4.16.1 4.16.2 4.16.3 4.16.3 Synchr 4.17.1 4.17.2 4.17.2 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.6 4.16.2.7 4.16.2.8 Math arran 4.16.3.1 4.16.3.2 4.16.3.3 conization a Barriers a device_ev | Asynchronous errors with a secondary queue | 2866 2877 2922 2922 2922 2933 3100 3111 312 312 312 312 312 312 312 312 31 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 4.16.3 4.16.3 Synchr 4.17.1 4.17.2 4.17.3 4.17.3 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.6 4.16.2.7 4.16.2.8 Math arra 4.16.3.1 4.16.3.2 4.16.3.3 conization a Barriers a device_ev | Asynchronous errors with a secondary queue class interface in types es es Vec interface Aliases Swizzles Swizzles Swizzled vec class Rounding modes Memory layout and alignment Considerations for endianness Performance note ' types Math array interface Aliases Memory layout and alignment in types erformance note ' types Math array interface Aliases Memory layout and alignment ind atomics ind fences ent class erences | 2866 2877 2922 2922 2923 3100 3111 312 312 312 312 312 312 312 312 31 |
| 4.16 | 4.15.2 Data ty 4.16.1 4.16.2 4.16.3 Synchr 4.17.1 4.17.2 4.17.3 4.17.4 | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.6 4.16.2.7 4.16.2.8 Math arra 4.16.3.1 4.16.3.2 4.16.3.3 conization a Barriers a device_ev Atomic re | Asynchronous errors with a secondary queue | 28662 28772 29222 29222 2933310 310 3111 3122312 31223 31223323 32332323 3244325 325235 325235 325235 325235 325235 325255 325255 325255 3252555 32525555 3252555555 325255555555 |
| 4.164.174.18 | 4.15.2 Data ty 4.16.1 4.16.2 4.16.3 Synchr 4.17.1 4.17.2 4.17.3 4.17.4 Stream | 4.15.1.4 Exception /pes Scalar dat Vector typ 4.16.2.1 4.16.2.2 4.16.2.3 4.16.2.4 4.16.2.5 4.16.2.6 4.16.2.7 4.16.2.8 Math arra 4.16.3.1 4.16.3.2 4.16.3.3 conization a Barriers a device_ev Atomic re Atomic typ | Asynchronous errors with a secondary queue | 2866 2877 2922 2922 2933 310 3111 3111 312 312 312 312 312 312 312 3 |

| | | 4.18.2 Synchronization | |
|---|--|---|--|
| | | 4.18.3 Implicit flush | |
| | | 4.18.4 Performance note | |
| | 4.19 | SYCL built-in functions for SYCL host and device | |
| | | 4.19.1 Description of the built-in types available for SYCL host and device | |
| | | 4.19.2 Work-item functions | |
| | | 4.19.3 Function objects | |
| | | 4.19.4 Algorithms library | |
| | | 4.19.4.1 any of, all of and none of | |
| | | 4.19.4.2 reduce | |
| | | 4.19.4.3 exclusive scan and inclusive scan | |
| | | 4 19 5 Group functions 357 | |
| | | 4 19 5 1 group broadcast 358 | |
| | | 4 1952 group harrier 358 | |
| | | 4 19 5 3 group any of group all of and group none of | |
| | | 4.19.5.4 group roduce 360 | |
| | | 4.19.5.4 group avaluative scan and group inclusive scan | |
| | | 4.19.5.5 group_exclusive_scall and group_inclusive_scall | |
| | | 4.19.0 Main functions | |
| | | 4.19.7 Integet functions | |
| | | 4.19.8 Common functions | |
| | | 4.19.9 Geometric functions | |
| | | 4.19.10 Relational functions | |
| | | 4.19.11 vector data load and store functions | |
| | | 4.19.12 Synchronization functions | |
| | | 4.19.13 printf function | |
| | | - | |
| 5 | SVC | T. Device Compiler 375 | |
| 5 | SYC 5 1 | CL Device Compiler 375 | |
| 5 | SYC 5.1 5.2 | CL Device Compiler 375 Offline compilation of SYCL source files 375 Naming of kernels 375 | |
| 5 | SYC 5.1 5.2 5.3 | CL Device Compiler 375 Offline compilation of SYCL source files 375 Naming of kernels 375 Compilation of functions 376 | |
| 5 | SYC 5.1 5.2 5.3 5.4 | CL Device Compiler 375 Offline compilation of SYCL source files 375 Naming of kernels 375 Compilation of functions 376 Language restrictions for device functions 376 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 | CL Device Compiler 375 Offline compilation of SYCL source files 375 Naming of kernels 375 Compilation of functions 376 Language restrictions for device functions 376 Built in scalar data types 377 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 | CL Device Compiler 375 Offline compilation of SYCL source files 375 Naming of kernels 375 Compilation of functions 376 Language restrictions for device functions 376 Built-in scalar data types 377 Preprocessor directives and macros 378 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 | CL Device Compiler 375 Offline compilation of SYCL source files 375 Naming of kernels 375 Compilation of functions 376 Language restrictions for device functions 376 Built-in scalar data types 377 Preprocessor directives and macros 378 Karnel attributes 370 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 | CL Device Compiler 375 Offline compilation of SYCL source files 375 Naming of kernels 375 Compilation of functions 376 Language restrictions for device functions 376 Built-in scalar data types 377 Preprocessor directives and macros 378 Kernel attributes 379 5.7.1 Corre kernel attributes 370 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 | St. Device Compiler 375 Offline compilation of SYCL source files 375 Naming of kernels 375 Compilation of functions 376 Language restrictions for device functions 376 Built-in scalar data types 377 Preprocessor directives and macros 378 Kernel attributes 379 5.7.1 Core kernel attributes 379 5.7.2 Example attribute suntar 381 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 | ST. Device Compiler 375 Offline compilation of SYCL source files 375 Naming of kernels 375 Compilation of functions 376 Language restrictions for device functions 376 Built-in scalar data types 377 Preprocessor directives and macros 378 Kernel attributes 379 5.7.1 Core kernel attributes 379 5.7.2 Example attribute syntax 381 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 | ST. Device Compiler 375 Offline compilation of SYCL source files 375 Naming of kernels 375 Compilation of functions 376 Language restrictions for device functions 376 Built-in scalar data types 377 Preprocessor directives and macros 378 Kernel attributes 379 5.7.1 Core kernel attributes 379 5.7.2 Example attribute syntax 381 5.7.3 Deprecated attribute syntax 381 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 | XL Device Compiler 375 Offline compilation of SYCL source files 375 Naming of kernels 375 Compilation of functions 376 Language restrictions for device functions 376 Built-in scalar data types 377 Preprocessor directives and macros 378 Kernel attributes 379 5.7.1 Core kernel attributes 379 5.7.2 Example attribute syntax 381 5.7.3 Deprecated attribute syntax 381 Address-space deduction 382 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 | XL Device Compiler 375 Offline compilation of SYCL source files 375 Naming of kernels 375 Compilation of functions 376 Language restrictions for device functions 376 Built-in scalar data types 377 Preprocessor directives and macros 378 Kernel attributes 379 5.7.1 Core kernel attributes 379 5.7.2 Example attribute syntax 381 5.7.3 Deprecated attribute syntax 381 5.8.1 Address space assignment 382 5.8.1 Address space assignment 382 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 | ZL Device Compiler375Offline compilation of SYCL source files375Naming of kernels375Compilation of functions376Language restrictions for device functions376Built-in scalar data types377Preprocessor directives and macros378Kernel attributes3795.7.1Core kernel attributes3795.7.2Example attribute syntax3815.7.3Deprecated attribute syntax3815.8.1Address space deduction3825.8.2Common address space deduction rules3825.8.2Common address space deduction rules3825.8.2Common address space deduction rules3825.8.2Common address space deduction rules3825.8.2Common address space deduction rules3825.8.3Comparison address space deduction rules3825.8.4Comparison address space deduction rules3825.8.2Common address space deduction rules3825.8.3Comparison address space deduction rules3825.8.4Comparison address space deduction rules3825.8.5Comparison address space deduction rules3825.8.1Comparison address space deduction rules3825.8.2Comparison address space deduction rules3825.8.3Comparison address space deduction rules3825.8.4Comparison address space deduction rules3825.8.5Comparison address space deduction rules3 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 | ZL Device Compiler375Offline compilation of SYCL source files375Naming of kernels375Compilation of functions376Language restrictions for device functions376Built-in scalar data types377Preprocessor directives and macros378Kernel attributes3795.7.1Core kernel attributes3795.7.2Example attribute syntax3815.7.3Deprecated attribute syntax3815.8.1Address space deduction3825.8.2Common address space3825.8.3Generic as default address space3825.8.4L fore her her her her her her her her her h | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 | ZL Device Compiler375Offline compilation of SYCL source files375Naming of kernels375Compilation of functions376Language restrictions for device functions376Built-in scalar data types377Preprocessor directives and macros378Kernel attributes3795.7.1Core kernel attributes5.7.2Example attribute syntax3815.7.3Deprecated attribute syntax3825.8.1Address space deduction3825.8.2Common address space3835.8.4Inferred address space384SWCHWe filled | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 | L Device Compiler375Offline compilation of SYCL source files375Naming of kernels375Compilation of functions376Language restrictions for device functions376Built-in scalar data types377Preprocessor directives and macros378Kernel attributes3795.7.1Core kernel attributes5.7.2Example attribute syntax3783815.7.3Deprecated attribute syntax381381Address-space deduction3825.8.1Address space assignment3823835.8.3Generic as default address space383384SYCL offline linking384Componentic of the function of the functi | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 | JL Device Compiler375Offline compilation of SYCL source files375Naming of kernels375Compilation of functions376Language restrictions for device functions376Built-in scalar data types377Preprocessor directives and macros378Kernel attributes3795.7.1Core kernel attributes5.7.2Example attribute syntax3813815.7.3Deprecated attribute syntax3823825.8.1Address space deduction rules3833845.8.4Inferred address space3845.9.1SYCL offline linking3845.9.1SYCL functions and member functions linkage | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.8 | CL Device Compiler375Offline compilation of SYCL source files375Naming of kernels375Compilation of functions376Language restrictions for device functions376Built-in scalar data types377Preprocessor directives and macros378Kernel attributes3795.7.1Core kernel attributes3795.7.2Example attribute syntax3815.7.3Deprecated attribute syntax381Address-space deduction3825.8.1Address space designment3825.8.2Common address space3835.8.4Inferred address space383SYCL offline linking3845.9.1SYCL functions and member functions linkage384Y. Extensions385 | |
| 6 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.8 5.9 SYC 6.1 | CL Device Compiler375Offline compilation of SYCL source files375Naming of kernels375Compilation of functions376Language restrictions for device functions376Built-in scalar data types377Preprocessor directives and macros378Kernel attributes3795.7.1Core kernel attributes3795.7.2Example attribute syntax3815.7.3Deprecated attribute syntax381Address-space deduction3825.8.1Address space deduction rules3825.8.2Common address space3835.8.4Inferred address space383SYCL offline linking3845.9.1SYCL functions and member functions linkage385Definition of an extension385385 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.8 5.9 SYC 6.1 6.2 | ZL Device Compiler375Offline compilation of SYCL source files375Naming of kernels375Compilation of functions376Language restrictions for device functions376Built-in scalar data types377Preprocessor directives and macros378Kernel attributes3795.7.1Core kernel attributes5.7.2Example attributes syntax373Deprecated attribute syntax374381Address-space deduction3825.8.1Address space assignment3823835.8.2Common address space3833845.9.1SYCL offline linking3845.9.1SYCL functions and member functions linkage3822Definition of an extension384385Definition of an extension385Definition of an extension385386Definition of an extension387385Definition of an extension385386385387385388385389385380385381386 <t< th=""><th></th></t<> | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.9 SYC 6.1 6.2 6.3 | ZL Device Compiler375Offline compilation of SYCL source files375Naming of kernels375Compilation of functions376Language restrictions for device functions376Built-in scalar data types377Preprocessor directives and macros378Kernel attributes3795.7.1Core kernel attributes5.7.2Example attribute syntax37.3Deprecated attribute syntax3815.7.3Deprecated attribute syntax3823815.8.1Address space deduction3823825.8.3Generic as default address space3833845.9.1SYCL functions and member functions linkage3843845.9.1SYCL functions and member functions linkage385Predefined macros386385Definition of an extension385Dervice aspects and conditional features387387Saperts and conditional features388Saperts and conditional features389Saperts and conditional features380Saperts and conditional features381Saperts and conditional features382Saperts and conditional features384Saperts and conditional features385Saperts and conditional features386Saperts and conditional features387Saperts and conditional features387Saperts and conditional features387Saperts and conditional | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.9 SYC 6.1 6.2 6.3 6.4 | ZL Device Compiler375Offline compilation of SYCL source files375Naming of kernels375Compilation of functions376Language restrictions for device functions376Built-in scalar data types377Preprocessor directives and macros378Kernel attributes3795.7.1Core kernel attributes3795.7.2Example attribute syntax3815.7.3Deprecated attribute syntax3814ddress-space deduction3825.8.1Address space assignment3825.8.2Common address space deduction rules3835.8.4Inferred address space383SYCL offline linking3845.9.1SYCL functions and member functions linkage3845.9.1SYCL functions and member functions linkage386Definition of an extension385386Device aspects and conditional features387Packends387Packends387 | |
| 5 | SYC 5.1 5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.9 SYC 6.1 6.2 6.3 6.4 6.5 | ZL Device Compiler375Offline compilation of SYCL source files375Naming of kernels375Compilation of functions376Language restrictions for device functions376Built-in scalar data types377Preprocessor directives and macros378Kernel attributes3795.7.1Core kernel attributes3795.7.2Example attribute syntax3815.7.3Deprecated attribute syntax3814ddress-space deduction3825.8.1Address space assignment3825.8.2Common address space deduction rules3825.8.3Generic as default address space3835.8.4Inferred address space3835.8.5SyCL offline linking3845.9.1SYCL functions and member functions linkage3845.9.1SYCL functions and member functions linkage386Definition of an extension385386Device aspects and conditional features387Conditional features387 <td< td=""><td></td></td<> | |

| A | Info | rmation descriptors 3 | 91 |
|----|------------------|--|-------------------|
| | A.1 | Platform information descriptors | 9 1 |
| | A.2 | Context information descriptors | 91 |
| | A.3 | Device information descriptors | 91 |
| | A.4 | Queue information descriptors | 94 |
| | A.5 | Kernel information descriptors | 94 |
| | A.6 | Event information descriptors | 95 |
| D | Foot | tumo coto | 07 |
| D | R 1 | Full feature set | 107 |
| | \mathbf{B}_{1} | Paduced feature set | 207 |
| | D.2 B 3 | | 107 |
| | В.3 В.4 | | 107 |
| | D.7 | | 91 |
| С | Host | t backend specification 3 | 99 |
| | C.1 | Mapping of the SYCL programming model on the host | 99 |
| | | C.1.1 SYCL memory model on the host | 99 |
| | C.2 | Interoperability with the host application | .00 |
| D | One | nCL backend specification | 101 |
| ν | D 1 | SYCL for OpenCL framework | 101 |
| | D.2 | Mapping of SYCL programming model on top of OpenCL | 101 |
| | 2.2 | D.2.1 Platform mixed version support | 02 |
| | | D.2.2 OpenCL memory model | 02 |
| | | D.2.3 OpenCL resources managed by SYCL application | 02 |
| | D.3 | Interoperability with the OpenCL API | 103 |
| | D.4 | Programming interface 4 | 105 |
| | | D.4.1 Reference counting | 07 |
| | | D.4.2 Errors and limitations | 07 |
| | | D.4.3 Interoperability with modules | 07 |
| | | D.4.3.1 Free functions | -08 |
| | | D.4.4 Interoperability with kernels | r09 |
| | | D.4.5 OpenCL kernel conventions and SYCL | r <mark>09</mark> |
| | | D.4.6 Data types | 10 |
| | D.5 | Preprocessor directives and macros | -11 |
| | | D.5.1 Offline linking with OpenCL C libraries | 11 |
| | D.6 | SYCL support of non-core OpenCL features | 13 |
| | | D.6.1 Half precision floating-point | 13 |
| | | D.6.2 Writing to 3D image memory objects | -14 |
| | | D.6.3 Interoperability with OpenGL | 14 |
| Е | Wha | at has changed from previous versions 4 | 15 |
| | E.1 | What has changed from SYCL 1.2.1 to SYCL 2020 | 15 |
| De | forer | - | 120 |
| ĸe | ieren | 4 | :20 |
| Gl | ossar | y 4 | 23 |

CONTENTS

List of Tables

| 3.1 | Combined requirement from two different accessor access modes within the same command | |
|------|---|----|
| | group. The rules are commutative and associative | 39 |
| | | |
| 4.1 | Common special member functions for reference semantics | 61 |
| 4.2 | Common hidden friend functions for reference semantics | 62 |
| 4.3 | Common special member functions for by-value semantics | 63 |
| 4.4 | Common hidden friend functions for by-value semantics | 63 |
| 4.5 | Traits for properties | 65 |
| 4.6 | Common member functions of the SYCL property interface | 66 |
| 4.7 | Constructors of the SYCL property_list class | 66 |
| 4.8 | Standard device selectors included with all SYCL implementations | 67 |
| 4.9 | Constructors of the SYCL platform class | 69 |
| 4.9 | Constructors of the SYCL platform class | 70 |
| 4.10 | Member functions of the SYCL platform class | 70 |
| 4.10 | Member functions of the SYCL platform class | 71 |
| 4.11 | Static member functions of the SYCL platform class | 71 |
| 4.12 | Platform information descriptors | 71 |
| 4.13 | Constructors of the SYCL context class | 73 |
| 4.14 | Member functions of the context class | 73 |
| 4.14 | Member functions of the context class | 74 |
| 4.15 | Context information descriptors | 74 |
| 4.15 | Context information descriptors | 75 |
| 4.16 | Constructors of the SYCL device class | 77 |
| 4.17 | Member functions of the SYCL device class | 77 |
| 4.17 | Member functions of the SYCL device class | 78 |
| 4.17 | Member functions of the SYCL device class | 79 |
| 4.18 | Static member functions of the SYCL device class | 79 |
| 4.19 | Device information descriptors | 79 |
| 4.19 | Device information descriptors | 80 |
| 4.19 | Device information descriptors | 81 |
| 4.19 | Device information descriptors | 82 |
| 4.19 | Device information descriptors | 83 |
| 4.19 | Device information descriptors | 84 |
| 4.19 | Device information descriptors | 85 |
| 4.19 | Device information descriptors | 86 |
| 4.19 | Device information descriptors | 87 |
| 4.19 | Device information descriptors | 88 |
| 4 19 | Device information descriptors | 89 |
| 4.19 | Device information descriptors | 90 |
| 4.19 | Device information descriptors | 91 |
| 4.20 | Device aspects defined by the core SYCL specification | 93 |
| 4 20 | Device aspects defined by the core SYCL specification | 94 |
| 4 21 | Constructors of the queue class | 97 |
| 1.41 | constructors of the queue ends | 1 |

| 4.21 | Constructors of the queue class | 98 |
|------|---|-----|
| 4.21 | Constructors of the queue class | 99 |
| 4.22 | Member functions for queue class | 99 |
| 4.22 | Member functions for queue class | 100 |
| 4.22 | Member functions for queue class | 101 |
| 4.22 | Member functions for queue class | 102 |
| 4.22 | Member functions for queue class | 103 |
| 4.23 | Queue information descriptors | 103 |
| 4.24 | Properties supported by the SYCL queue class | 104 |
| 4.25 | Constructors of the gueue property classes | 104 |
| 4.26 | Constructors of the event class | 106 |
| 4.27 | Member functions for the event class | 106 |
| 4.27 | Member functions for the event class | 107 |
| 4.27 | Member functions for the event class | 108 |
| 4.28 | Event class information descriptors | 108 |
| 4.29 | Profiling information descriptors for the SYCL event class | 109 |
| 4.30 | SYCL Default Allocators | 110 |
| 4.31 | Constructors of the buffer class | 114 |
| 4.31 | Constructors of the buffer class | 115 |
| 4.31 | Constructors of the buffer class | 116 |
| 4.31 | Constructors of the buffer class | 117 |
| 4 31 | Constructors of the buffer class | 118 |
| 4 32 | Member functions for the buffer class | 118 |
| 4 32 | Member functions for the buffer class | 119 |
| 4 32 | Member functions for the buffer class | 120 |
| 4 32 | Member functions for the buffer class | 121 |
| 4 33 | Properties supported by the SYCL buffer class | 121 |
| 4 33 | Properties supported by the SYCL buffer class | 122 |
| 4.34 | Constructors of the buffer property classes | 122 |
| 4.35 | Member functions of the buffer property classes | 122 |
| 4.36 | Constructors of the unsampled image class template | 127 |
| 4.36 | Constructors of the unsampled image class template | 128 |
| 4.36 | Constructors of the unsampled_image class template | 129 |
| 4.36 | Constructors of the unsampled_image class template | 130 |
| 4.36 | Constructors of the unsampled_image class template | 131 |
| 4.36 | Constructors of the unsampled_image class template | 132 |
| 4.37 | Member functions of the unsampled_image class template | 133 |
| 4.38 | Constructors of the sampled_image class template | 135 |
| 4.38 | Constructors of the sampled_image class template | 136 |
| 4.38 | Constructors of the sampled_image class template | 137 |
| 4.39 | Member functions of the sampled_image class template | 137 |
| 4.39 | Member functions of the sampled_image class template | 138 |
| 4.40 | Properties supported by the SYCL image classes | 138 |
| 4.40 | Properties supported by the SYCL image classes | 139 |
| 4.41 | Constructors of the image property classes | 139 |
| 4.42 | Member functions of the image property classes | 139 |
| 4.43 | Enumeration of access modes available to accessors | 143 |
| 4.44 | Enumeration of access modes available to accessors | 144 |
| 4.45 | Enumeration of access tags available to accessors | 145 |
| 4.46 | Description of all the device buffer accessor capabilities | 148 |
| 4.47 | Member types of the accessor class template buffer specialization | 152 |

| 4.48 | Constructors of the accessor class template buffer specialization | 53 |
|--|--|--|
| 4.48 | Constructors of the accessor class template buffer specialization | 54 |
| 4.48 | Constructors of the accessor class template buffer specialization | 55 |
| 4.49 | Member functions of the accessor class template buffer specialization | 55 |
| 4.49 | Member functions of the accessor class template buffer specialization | 56 |
| 4.49 | Member functions of the accessor class template buffer specialization | 57 |
| 4.49 | Member functions of the accessor class template buffer specialization | 58 |
| 4.50 | Properties supported by the SYCL accessor class | 58 |
| 4.51 | Constructors of the accessor property classes | 58 |
| 4.52 | Member types of the host accessor class template | 52 |
| 4.53 | Constructors of the host accessor class template | 53 |
| 4.53 | Constructors of the host accessor class template | 54 |
| 4.53 | Constructors of the host accessor class template | 55 |
| 4 54 | Member functions of the host accessor class template | 55 |
| 4 54 | Member functions of the host_accessor class template 11 | 56 |
| 4 54 | Member functions of the host_accessor class template | 57 |
| 4 55 | Description of all the local accessor canabilities | 58 |
| 4 56 | Member types of the accessor class template local specialization | 70 |
| 4.50 | Constructors of the accessor class template local specialization | 70 |
| 4.57 | Constructors of the accessor class template local specialization 17 | 71 |
| 4.57 | Member functions of the accessor class template local specialization 17 | 71 |
| 4.50 | Member functions of the accessor class template local specialization | 72 |
| 4.50 | Member functions of the accessor class template local specialization | 12 |
| 4.50 | Description of all the image accessor capabilities | · 5 7 / |
| 4.59 | Constructors of the accessor capabilities | 14 16 |
| 4.00 | Mamber functions of the accessor class template image specialization | 16 |
| | | |
| 4.01 | Member functions of the accessor class template image specialization | 10 |
| 4.61 | Member functions of the accessor class template image specialization | 17 17 |
| 4.61 4.62 | Member functions of the accessor class template image specialization | 77 34 |
| 4.61 4.62 4.62 4.62 | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Constructors of the SYCL multi_ptr class template 18 Operators of multi_netr class 18 Operators of multi_netr class 18 | 77 34 35 |
| 4.61 4.62 4.62 4.63 | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Operators of multi_ptr class 18 Operators of multi_ptr class 18 | 77 34 35 35 |
| 4.61 4.62 4.62 4.63 4.63 | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 | 77 34 35 35 36 |
| 4.61 4.62 4.62 4.63 4.63 4.63 | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of multi_ptr class 18 Operators of multi_ptr | 77 34 35 35 36 37 |
| 4.61 4.62 4.62 4.63 4.63 4.63 4.63 | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Uperators of multi_ptr class 18 Uperators of multi_ptr class 18 Uperators of multi_ptr class 18 Member functions of functions of the multi_ptr class 18 Member functions of functions of the multi_ptr class 18 | 77 34 35 35 36 37 37 |
| 4.61 4.62 4.62 4.63 4.63 4.63 4.63 4.64 4.65 | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Constructors of multi_ptr class 18 Operators of multi_ptr class 18 Member functions of the multi_ptr class 18 Member functions of the multi_ptr class 18 Member functions of the multi_ptr class 18 Midden friend functions of the multi_ptr class 18 | 77 34 35 35 36 37 37 38 |
| 4.61 4.62 4.62 4.63 4.63 4.63 4.63 4.64 4.65 4.65 | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Operators of multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 | 77 34 35 35 36 37 37 38 39 |
| $\begin{array}{r} 4.61\\ 4.62\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.64\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\end{array}$ | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of multi_ptr class 18 Medden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Addressing modes description 19 State in the intervent 19 | 77 34 35 35 36 37 38 39 55 |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.64\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.66\end{array}$ | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Addressing modes description 19 Filtering modes description 19 Filtering modes description 19 Operators of whether the label of t | 77 34 35 36 37 38 39 55 50 50 |
| $\begin{array}{r} 4.61\\ 4.61\\ 4.62\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.68\\ 4.69\end{array}$ | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Addressing modes description 19 Filtering modes description 19 Coordinate normalization modes description 19 Condinate normalization modes description 19 | 77 34 35 36 37 37 38 39 5 5 5 5 5 |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.68\\ 4.69\\$ | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of multi_ptr class 18 Member functions of multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Addressing modes description 19 Filtering modes description 19 Constructors the sampler class 19 Member functions of the multi_ptr class 19 State and the sampler class 19 Coordinate normalization modes description 19 Constructors the sampler class 19< | 77 34 35 36 37 38 39 55 50 50 50 50 50 50 50 50 50 50 50 50 |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.68\\ 4.69\\ 4.70\\ 4.70\\ 4.71\\ 4.68\\ 4.69\\ 4.70\\ 4.71\\ 4.68\\ 4.70\\ 4.71\\ 4.71\\ 4.72\\ 4.70\\ 4.71\\ 4.72\\ 4.70\\ 4.72\\$ | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of multi_ptr class 18 Member functions of multi_ptr class 18 Mether functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Addressing modes description 19 Filtering modes description 19 Coordinate normalization modes description 19 Member functions for the sampler class 19 Member functions for the sampler class 19 Constructors the sampler class 19 Member functions for the sampler class 19 Member functions for the sampler class 19 Member functions | 77 34 35 36 37 37 38 39 55 56 16 16 16 16 16 16 16 16 16 16 16 16 16 |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.68\\ 4.69\\ 4.70\\ 4.71\\ 4.71\end{array}$ | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of multi_ptr class 18 Member functions of multi_ptr class 18 Member functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 19 Filtering modes description 19 Coordinate normalization modes description 19 Constructors the sampler class 19 Member functions for the sampler class 19 Type of USM allocations 19 Constructors the sampler class 19 Member functions for the sampler class 19 Constructors the sampler class 19 Member functions for the sampler clas | 77 34 35 36 37 7 8 39 55 56 16 18 10 10 10 10 10 10 10 10 10 10 10 10 10 |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.68\\ 4.69\\ 4.70\\ 4.71\\ 4.72\\ 4.72\end{array}$ | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of multi_ptr class 18 Member functions of multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Addressing modes description 19 Filtering modes description 19 Constructors the sampler class 19 Member functions for the sampler class 19 Member functions of the differe | 77 34 35 36 37 38 39 55 56 10 10 10 10 10 10 10 10 10 10 10 10 10 |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.68\\ 4.69\\ 4.70\\ 4.71\\ 4.72\\ 4.73\end{array}$ | Member functions of the accessor class template image specialization 17 Member functions of the SYCL multi_ptr class template 18 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of the multi_ptr class 18 Member functions of the multi_ptr class 18 Member functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Addressing modes description 19 Filtering modes description 19 Coordinate normalization modes description 19 Constructors the sampler class 19 Member functions for the sampler class 19 Member functions of the different kinds of USM allocation 19 Constructors the sampler class 19 Member functions for the sampler class 19 Member funct | 77 34 35 36 37 37 38 39 55 55 56 18 18 18 18 18 18 18 18 18 18 18 18 18 |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.68\\ 4.69\\ 4.70\\ 4.71\\ 4.72\\ 4.73\\ 4.73\end{array}$ | Member functions of the accessor class template image specialization 17 Member functions of the SYCL multi_ptr class template 18 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 19 Filtering modes description 19 Coordinate normalization modes description 19 Constructors the sampler class 19 Member functions for the sampler class 19 Member functions for the sampler class 19 Characteristics of the different kinds of USM allocation 19 Characteristics of the different kinds of USM allocation 19 < | 77 34 35 36 37 38 39 55 36 37 38 39 35 36 37 38 39 35 36 37 38 39 35 36 37 38 39 35 36 37 38 39 35 36 37 38 39 35 36 <td< td=""></td<> |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.68\\ 4.69\\ 4.70\\ 4.71\\ 4.72\\ 4.73\\ 4.74\\ 4.74\end{array}$ | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of multi_ptr class 18 Member functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Addressing modes description 19 Filtering modes description 19 Constructors the sampler class 19 Member functions for the sampler class 19 Member functions for the sampler class 19 Constructors the sampler class 19 Characteristics of the different kinds of USM allocation 19 Summary of types used to identify points in an index space, and ranges over which those points 21 Constructors of the range class template </td <td>77 34 35 36 37 38 39 55 36 37 38 39 55 36 37 38 39 55 36 37 38 39 35 36 37 38 39 35 36 37 38 39 35 36 37 38 39 35 36 36 3 4 5 36 34 3 4 5 36 3 4 5 36 36 3 4 5 36 3 4 5 36</td> | 77 34 35 36 37 38 39 55 36 37 38 39 55 36 37 38 39 55 36 37 38 39 35 36 37 38 39 35 36 37 38 39 35 36 37 38 39 35 36 36 3 4 5 36 34 3 4 5 36 3 4 5 36 36 3 4 5 36 3 4 5 36 |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.68\\ 4.69\\ 4.70\\ 4.71\\ 4.72\\ 4.73\\ 4.74\\ 4.75\\ 4.74\\ 4.75\\ 4.74\\ 4.75\\ 4.74\\ 4.75\\ 4.74\\ 4.75\\$ | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of multi_ptr class 18 Member functions of multi_ptr class 18 Member functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Addressing modes description 19 Filtering modes description 19 Constructors the sampler class 19 Member functions for the sampler class 19 Member functions for the sampler class 19 Characteristics of the different kinds of USM allocation 19 Summary of types used to identify points in an index space, and ranges over which those points 21 Constructors of the range class templ | 77 34 35 35 36 37 37 38 39 35 30 37 38 39 39 35 34 36 35 36 36 37 38 39 39 35 30 36 31 37 32 36 33 37 39 35 30 36 31 37 32 37 33 39 34 36 37 36 38 39 39 30 34 36 34 36 |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.68\\ 4.69\\ 4.70\\ 4.71\\ 4.72\\ 4.73\\ 4.74\\ 4.75\\ 4.76\\ 1.76\\$ | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of multi_ptr class 18 Member functions of multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Addressing modes description 19 Filtering modes description 19 Constructors the sampler class 19 Member functions for the sampler class 19 Member functions for the sampler class 19 Characteristics of the different kinds of USM allocation 19 Summary of types used to identify points in an index space, and ranges over which those points 21 Constructors of the range class template 21 Member functions of the range class template 21 Member functions of | 77 34 35 36 37 38 39 55 56 37 38 39 55 56 37 38 39 55 56 37 38 39 55 36 37 38 39 55 36 37 37 38 39 55 36 37 37 38 39 55 36 37 37 38 35 35 36 37 37 38 35 35 36 37 37 38 35 35 36 37 37 38 35 35 36 37 37 38 35 35 36 37 37 38 35 35 36 37 37 38 35 35 36 37 37 38 35 35 36 37 37 38 39 55 36 37 37 38 39 55 36 37 37 38 39 55 36 37 37 38 39 55 36 37 37 38 39 55 36 37 37 38 39 55 36 37 37 38 39 55 36 37 37 38 39 35 36 37 37 38 39 35 37 37 38 39 35 36 37 37 38 39 35 36 37 37 38 39 35 35 36 37 37 38 39 35 37 37 38 39 35 35 35 36 37 37 38 39 35 35 36 37 37 38 39 35 35 37 37 38 39 35 35 35 35 35 35 35 35 35 35 35 35 35 |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.68\\ 4.69\\ 4.70\\ 4.71\\ 4.72\\ 4.73\\ 4.74\\ 4.75\\ 4.76\\ 4.76\\ 4.76\\ 4.76\end{array}$ | Member functions of the accessor class template image specialization 17 Member functions of the accessor class template image specialization 17 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Addressing modes description 19 Filtering modes description 19 Constructors the sampler class 19 Member functions for the sampler class 19 Member functions of the different kinds of USM allocation 19 Constructors of the different kinds of USM allocation 19 Constructors of the range class template 21 Member functions of the range class template 21 Member functions of the range class template 21 Constructors of the range class template 21 <td>77 34 35 36 37 38 9 5 5 6 12<</td> | 77 34 35 36 37 38 9 5 5 6 12< |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.68\\ 4.69\\ 4.70\\ 4.71\\ 4.72\\ 4.73\\ 4.74\\ 4.75\\ 4.76\\ 4.76\\ 4.76\\ 4.77\end{array}$ | Member functions of the accessor class template image specialization 17 Member functions of the SYCL multi_ptr class template 18 Constructors of the SYCL multi_ptr class template 18 Operators of multi_ptr class 18 Member functions of multi_ptr class 18 Member functions of the multi_ptr class 18 Hidden friend functions of the multi_ptr class 18 Addressing modes description 19 Filtering modes description 19 Constructors the sampler class 19 Member functions for the sampler class 19 Member functions of the different kinds of USM allocation 19 Constructors of the different kinds of USM allocation 19 Constructors of the range class template 21 | 77 34 5 35 6 37 7 8 9 9 5 9 6 9 8 9 8 13 14 15 16 7 |
| $\begin{array}{c} 4.61\\ 4.61\\ 4.62\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.63\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.65\\ 4.66\\ 4.67\\ 4.70\\ 4.71\\ 4.72\\ 4.73\\ 4.74\\ 4.75\\ 4.76\\ 4.76\\ 4.76\\ 4.77\\ 4.78\end{array}$ | Member functions of the accessor class template image specialization17Member functions of the SYCL multi_ptr class template18Constructors of the SYCL multi_ptr class template18Operators of multi_ptr class18Operators of multi_ptr class18Operators of multi_ptr class18Operators of multi_ptr class18Member functions of multi_ptr class18Member functions of multi_ptr class18Member functions of the multi_ptr class18Member functions of the multi_ptr class18Hidden friend functions of the multi_ptr class18Addressing modes description19Filtering modes description19Constructors the sampler class19Member functions for the sampler class19Member functions of the different kinds of USM allocation19Constructors of the different kinds of USM allocation19Constructors of the arage class template21Member functions of the range class template21Member functions of the syCL range class template21Member functions of the syCL range class template21Member functions of the different kinds of USM allocation< | 77 34 5 35 6 7 7 34 5 35 6 7 7 34 5 35 6 7 7 34 5 35 6 7 7 34 5 35 6 7 7 34 5 35 6 7 7 7 6 10 10 10 10 10 10 10 10 10 10 10 10 10 |

| 4.80 | Member functions of the id class template | . 219 |
|-------|---|-------|
| 4.81 | Hidden friend functions of the id class template | . 219 |
| 4.81 | Hidden friend functions of the id class template | . 220 |
| 4.82 | Member functions for the item class | . 221 |
| 4.83 | Member functions for the nd item class | 223 |
| 4.83 | Member functions for the nd_item class | 224 |
| 4 83 | Member functions for the nd_item class | 225 |
| 4 84 | Member functions for the h item class | 226 |
| 4 84 | Member functions for the h_item class | 220 |
| 4 84 | Member functions for the h_item class | 227 |
| 4 85 | Member functions for the group class | 220 |
| 1.05 | Member functions for the group class | 230 |
| 4.05 | Member functions for the group class | . 231 |
| 4.05 | Member functions for the group class | . 232 |
| 4.05 | Member functions for the sub-group class | . 200 |
| 4.00 | Overlands of the reduction interface | . 234 |
| 4.8/ | | . 237 |
| 4.8/ | | . 238 |
| 4.8/ | Overloads of the reduction interface | . 239 |
| 4.88 | Constructors of the reducer class | . 241 |
| 4.88 | Constructors of the reducer class | . 242 |
| 4.89 | Member functions of the reducer class | . 242 |
| 4.89 | Member functions of the reducer class | . 243 |
| 4.90 | Operators of the reducer class | . 243 |
| 4.91 | Constructors of the handler class | . 247 |
| 4.92 | Member functions of the handler class | . 248 |
| 4.93 | Member functions of the handler class | . 248 |
| 4.93 | Member functions of the handler class | . 249 |
| 4.93 | Member functions of the handler class | . 250 |
| 4.93 | Member functions of the handler class | . 251 |
| 4.93 | Member functions of the handler class | . 252 |
| 4.94 | Constructor of the private_memory class | . 257 |
| 4.95 | Member functions of the private_memory class | . 258 |
| 4.96 | Member functions of the handler class | . 261 |
| 4.97 | Member functions of the kernel class | . 268 |
| 4.97 | Member functions of the kernel class | . 269 |
| 4.98 | Kernel class information descriptors | . 269 |
| 4.99 | Device-specific kernel information descriptors | . 270 |
| 4.100 | OKernel work-group information descriptors | . 271 |
| 4.10 | Member functions of the SYCL exception class | . 289 |
| 4.101 | 1 Member functions of the SYCL exception class | 290 |
| 4.102 | 2Member functions of the exception list | 290 |
| 4 102 | 2Member functions of the exception list | 291 |
| 4 103 | 3Values of the SYCL errc enum | 291 |
| 4 104 | 4SYCL error code helper functions | 291 |
| 4 104 | 5 Additional scalar data types supported by SYCI | 202 |
| 4 10 | 6Constructors of the SVCL up class template | 296 |
| 4 104 | 6Constructors of the SVCL upc class template | · 290 |
| 4.100 | 7Member functions for the SVCL use class template | · 271 |
| 4.10 | 7 Member functions for the SVCL use class template | · 271 |
| 4.10 | Momber functions for the SYCL use class template | . 298 |
| 4.10 | / Wiember functions for the SYCL vec class template | . 299 |
| 4.10 | / Member functions for the SYCL vec class template | . 500 |

| 4.107 Member functions for the SYCL vec class template |
|---|
| 4.108Hidden friend functions of the vec class template |
| 4.108Hidden friend functions of the vec class template |
| 4.108Hidden friend functions of the vec class template |
| 4.108Hidden friend functions of the vec class template |
| 4.108Hidden friend functions of the vec class template |
| 4.108Hidden friend functions of the vec class template |
| 4.108Hidden friend functions of the vec class template |
| 4.108Hidden friend functions of the vec class template |
| 4.108Hidden friend functions of the vec class template |
| 4.108Hidden friend functions of the vec class template |
| 4.109Rounding modes for the SYCL vec class template |
| 4.110Constructors of the SYCL marray class template |
| 4.111 Member functions for the SYCL marray class template |
| 4.111 Member functions for the SYCL marray class template 316 |
| 4.112Non-member functions of the marray class template 316 |
| 4 112Non-member functions of the marray class template 317 |
| 4.112Non-member functions of the marray class template |
| 4.112Non-member functions of the marray class template |
| 4.112Non-member functions of the marray class template |
| 4.112Non member functions of the marray class template |
| 4.112Non-member functions of the marray class template |
| 4.112Non-member functions of the marray class template |
| 4.112 Non-memoer functions of the SVCL device, event class |
| 4.113 Member functions of the device_event class |
| 4.114 Constructors of the SVCL storrig ref class templete |
| 4.115 Constructions of the STCL atomic_ref class template |
| 4.116 Member functions available on any object of type atomic_ref<1> |
| 4.110 Member functions available on any object of type atomic_rer<1> |
| 4.117 Additional member functions available on an object of type atomic_rel<1> for integral T |
| 4.117 Additional member functions available on an object of type atomic_rel<1> for floating point T |
| 4.118 Additional member functions available on an object of type atomic_rei<1> for floating-point T |
| 4.110 Additional member functions available on an object of type atomic_ret<1> 101 floating-point 1 |
| 4.119Additional member functions available on an object of type atomic_ref<1^> |
| 4.120 Constructors of the c1::syc1::atomic class template |
| 4.121 Member functions available on an object of type cl::sycl::atomic<1> |
| 4.121 Member functions available on an object of type cl::sycl::atomic<1> |
| 4.121 Member functions available on an object of type cl::sycl::atomic<1> |
| 4.121 Member functions available on an object of type c1::syc1::atomic<1> |
| 4.122 Global functions available on atomic types |
| $4.122 \text{Global functions available on atomic types} \dots \dots$ |
| 4.123 Operand types supported by the stream class |
| 4.124 Manipulators supported by the stream class |
| 4.125Constructors of the stream class |
| 4.126 Member functions of the stream class |
| 4.12/Global functions of the stream class |
| 4.128 Generic type name description, which serves as a description for all valid types of parameters to |
| kernet runctions [1] 348 4 129 Complete transmission 121 transmission |
| 4.128 Generic type name description, which serves as a description for all valid types of parameters to |
| kernet functions [1] |
| 4.128 Generic type name description, which serves as a description for all valid types of parameters to |
| kernel functions [1] |

| 4.129Member functions for the plus function object |
|---|
| 4.130Member functions for the multiplies function object |
| 4.131 Member functions for the bit_and function object |
| 4.132Member functions for the bit_or function object |
| 4.133Member functions for the bit_xor function object |
| 4.134Member functions for the logical_and function object |
| 4.135Member functions for the logical or function object |
| 4.136Member functions for the minimum function object |
| 4.137 Member functions for the maximum function object |
| 4.138Overloads for the any of function 354 |
| 4.139Overloads for the all of function |
| 4.140Overloads for the none of function |
| 4.141Overloads of the reduce function 355 |
| 4 142Overloads of the exclusive scan function 356 |
| 4 143Overloads of the inclusive scan function 357 |
| 4 144Overloads of the group broadcast function 358 |
| 4 145Overloads for the group harrier function 359 |
| 4 146Overloads for the group any of function 359 |
| 4.147Overloads for the group all of function 359 |
| 4.147 Overloads for the group all of function 360 |
| 4.148 Overloads for the group none of function 360 |
| 4 149Overloads of the group reduce function 360 |
| 4 150 Overloads of the group exclusive scan function 361 |
| 4.151Overloads of the group inclusive scan function 362 |
| 4.152Math functions which work on SVCL host and device. They correspond to Table 6.8 of the |
| OpenCL 1.2 specification [1] |
| |
| 4 152Math functions which work on SVCL host and device. They correspond to Table 6.8 of the |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 364 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 366 4.153Native math functions 366 4.153Native math functions 367 4.154Half precision math functions 367 4.155Integer functions which work on SYCL host and device, are available in the sycl namespace 368 4.156Common functions which work on SYCL host and device, are available in the sycl namespace 369 4.156Common functions which work on SYCL host and device, are available in the sycl namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1] 369 |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 364 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 366 4.153Native math functions 366 4.153Native math functions 367 4.154Half precision math functions 367 4.155Integer functions which work on SYCL host and device, are available in the syc1 namespace 368 4.155Integer functions which work on SYCL host and device, are available in the syc1 namespace 369 4.156Common functions which work on SYCL host and device, are available in the syc1 namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1] 369 4.156Common functions which work on SYCL host and device, are available in the syc1 namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1] 369 4.156Common functions which work on SYCL host and device, are available in the syc1 namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1] < |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 364 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.153Native math functions 366 4.153Native math functions 366 4.153Native math functions 367 4.154Half precision math functions 367 4.155Integer functions which work on SYCL host and device, are available in the syc1 namespace 368 4.155Integer functions which work on SYCL host and device, are available in the syc1 namespace 369 4.156Common functions which work on SYCL host and device, are available in the syc1 namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1] 369 4.156Common functions which work on SYCL host and device, are available in the syc1 namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1] 369 4.156Common functions which work on SYCL host and device, are available in the syc1 namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1] 369 |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 364 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 364 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.153Native math functions 366 4.153Native math functions 366 4.153Native math functions 367 4.154Half precision math functions 367 4.155Integer functions which work on SYCL host and device, are available in the syc1 namespace 368 4.155Integer functions which work on SYCL host and device, are available in the syc1 namespace 369 4.156Common functions which work on SYCL host and device, are available in the syc1 namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1] 370 4.157Geometric functions which work on SYCL host and device, are available in the syc1 namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1] 370 4.157Geometric functions which work on SYCL host and device, are available in the syc1 namespace. They correspond to Table 6.13 of the OpenCL 1.2 specification [1] 370 |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] |
| 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 364 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 365 4.152Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1] 366 4.153Mative math functions 366 4.153Native math functions 367 4.154Half precision math functions 367 4.155Integer functions which work on SYCL host and device, are available in the sycl namespace 368 4.155Integer functions which work on SYCL host and device, are available in the sycl namespace 369 4.156Common functions which work on SYCL host and device, are available in the sycl namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1] 370 4.157Geometric functions which work on SYCL host and device, are available in the sycl namespace. They correspond to Table 6.13 of the OpenCL 1.2 specification [1] 370 4.157Geometric functions which work on SYCL host and device, are available in the sycl namespace. They correspond to Table 6.13 of the OpenCL 1.2 specification [1] 370 4.157Geometric functions which work on SYCL host and device, are available in the sycl namespace. They correspond to Table 6.13 |

| 4.159 | PRelational functions for scalar data types and marray template class template class which work on | |
|-------|--|---|
| | SYCL host and device, are available in the sycl namespace | 4 |
| 5.1 | Fundamental data types supported by SYCL | 3 |
| 5.2 | Attributes supported by the SYCL General programming interface |) |
| 5.2 | Attributes supported by the SYCL General programming interface | 1 |
| C.1 | Mapping of SYCL memory regions into host memory regions | 9 |
| D.1 | Mapping of SYCL memory regions into OpenCL memory regions | 2 |
| D.2 | List of native types per SYCL object in the OpenCL backend | 4 |
| D.3 | List of native types per SYCL object on kernel code | 4 |
| D.6 | Example range mapping from SYCL enqueued three dimensional global range to OpenCL and | |
| | SYCL queries |) |
| D.7 | Scalar data type aliases supported by SYCL OpenCL backend | 1 |
| D.8 | SYCL support for OpenCL 1.2 extensions | 3 |

List of Figures

| 3.1 | Execution order of three command groups submitted to the same queue | 36 |
|-----|--|----|
| 3.2 | Execution order of three command groups submitted to the different queues | 36 |
| 3.3 | Actions performed when three command groups are submitted to two distinct queues, and poten- | |
| | tial implementation in an OpenCL SYCL backend by a SYCL runtime. Note that in this example, | |
| | each SYCL buffer $(b1, b2)$ is implemented as separate cl_mem objects per context | 38 |
| 3.4 | Requirements on overlapping vs non-overlapping sub-buffer | 40 |
| 3.5 | Execution of command groups when using host accessors | 40 |

LIST OF FIGURES

Listings

| code/anatomy.cpp | | |
|---|--|--|
| code/largesample.cpp | | |
| headers/backends.h | | |
| headers/paramTraits.h | | |
| headers/interop/typeTraitsBackendTraits.h | | |
| headers/interop/templateFunctionGetNative.h | | |
| headers/interop/templateFunctionMakeX.h | | |
| headers/common-reference.h | | |
| headers/common-byval.h | | |
| code/propertyExample.cpp | | |
| headers/properties.h | | |
| headers/deviceSelector.h | | |
| headers/platform.h | | |
| headers/context.h | | |
| headers/device.h | | |
| headers/deviceEnumClassAspect.h | | |
| headers/deviceBackendAspect.h | | |
| headers/queue.h | | |
| headers/event.h | | |
| headers/buffer.h | | |
| code/subbuffer.cpp | | |
| headers/unsampledImage.h | | |
| headers/sampledImage.h | | |
| headers/imageProperties.h | | |
| headers/accessTarget.h | | |
| headers/accessMode.h | | |
| headers/accessTags.h | | |
| 4.1 Accessor declaration | | |
| 4.2 Using an accessor to access a sub-range | | |
| 4.3 Device accessor class for buffers | | |
| headers/accessProperties.h | | |
| 4.4 Host accessor class for buffers | | |
| 4.5 Accessor class for locals | | |
| 4.6 Accessor interface for images | | |
| headers/multipointer.h | | |
| headers/multipointerlegacy.h | | |
| headers/pointer.h | | |
| headers/sampler.h | | |
| headers/range.h | | |
| headers/ndRange.h | | |
| headers/id.h | | |
| headers/item.h | | |
| headers/nditem.h | | |
| headers/hitem.h | | |

| headers/group.h | . 228 |
|--|-------|
| headers/subgroup.h | . 233 |
| code/reduction.cpp | . 235 |
| code/parallelreduce.cpp | . 235 |
| headers/reduction.h | . 237 |
| headers/reducer.h | . 239 |
| headers/commandGroupHandler.h | . 244 |
| headers/expressingParallelism/classKernelHandler.h | . 247 |
| headers/expressingParallelism/classKernelHandler/constructors.h | . 247 |
| headers/expressingParallelism/classKernelHandler/hasSpecializationConstant.h | . 247 |
| headers/expressingParallelism/classKernelHandler/getSpecializationConstant.h | . 247 |
| code/singletask.cpp | . 253 |
| code/singleTaskWithKernelHandler.cpp | . 253 |
| code/basicparallelfor.cpp | . 254 |
| code/basicParallelForItem.cop | . 254 |
| code/basicParallelForGeneric.cpp | 254 |
| code/basicParallelForIntegral.cpp | 254 |
| code/basicParallelForNumber cnp | 255 |
| code/narallelfor.cnn | 255 |
| code/parallelforharrier cnn | 256 |
| code/parallelEorWithKernelHandler.com | 256 |
| headers/priv h | 257 |
| code/parallelforworkgroup.com | 258 |
| code/parallelforworkgroup?cnp | 250 |
| code/parallelForWorkGroupWithKernelHandler.com | 259 |
| code/paranen of workoroup with Kernen randiel.epp | · 259 |
| headers/handler/useModule h | . 202 |
| headers/handler/heastracializationConstant h | . 202 |
| headers/handler/hastfragializationConstant.h | . 202 |
| headers/heatTeal/heatTeal/heatTeals/wears/heatTeals/heat | . 205 |
| headers/hostTask/hostTask/slossInterenHendle.h | . 205 |
| headers/hostTask/classInteropHandle.n | . 203 |
| headers/host lask/classifieropHandle/constructors.n | . 205 |
| headers/host lask/classInteropHandle/gethativeX.n. | . 205 |
| | . 200 |
| | . 267 |
| headers/module/modulesSynopsis.h | . 272 |
| headers/module/enumClassModuleState.h | . 274 |
| headers/module/class lemplateSpecializationId.h | . 275 |
| headers/module/class TemplateSpecializationId/constructors.h | . 275 |
| headers/module/class TemplateSpecializationId/specialmembers.h. | . 275 |
| headers/module/classTemplateModule.h | . 276 |
| headers/module/classTemplateModule/constructors.h | . 276 |
| headers/module/classTemplateModule/getContext.h | . 276 |
| headers/module/classTemplateModule/getDevices.h | . 276 |
| headers/module/classTemplateModule/hasKernel.h | . 276 |
| headers/module/classTemplateModule/getKernel.h | . 277 |
| headers/module/classTemplateModule/getKernelNames.h | . 277 |
| headers/module/classTemplateModule/isEmpty.h | . 277 |
| headers/module/classTemplateModule/beginAndEnd.h | . 277 |
| $headers/module/classTemplateModule/containsSpecializationConstants.h \ . \ . \ . \ . \ . \ . \ . \ . \ . \$ | . 277 |
| headers/module/classTemplateModule/nativeSpecializationConstant.h | . 277 |

| headers/module/classTemplateModule/hasSpecializationConstant.h |
|--|
| headers/module/classTemplateModule/setSpecializationConstant.h |
| headers/module/classTemplateModule/getSpecializationConstant.h |
| headers/module/freeFunctions.h |
| headers/module/namespaceThisModule/kernelName.h |
| headers/module/namespaceThisModule/freeFunctions.h |
| code/myfunctor.cpp |
| code/mykernel.cpp |
| code/mymodule.cpp |
| code/usingSpecConstants.cpp |
| code/handlingException.cpp |
| code/handlingErrorCode.cpp |
| code/handlingBackendErrorCode.cpp |
| headers/exception.h |
| headers/vec.h |
| headers/marray.h |
| headers/synchronization.h |
| headers/deviceEvent.h |
| headers/atomicref.h |
| headers/atomic.h |
| headers/atomicoperations.h |
| headers/stream.h |
| headers/functional.h |
| code/attributes.cpp |
| headers/extension/deviceKhr.h |
| headers/extension/deviceExtAcme.h |
| headers/extension/appMigration.h |
| headers/extension/aspectKhr.h |
| headers/extension/aspectVendor.h |
| headers/extension/backend.h |
| headers/platformInfo.h |
| headers/contextInfo.h |
| headers/deviceInfo.h |
| headers/queueInfo.h |
| headers/kernelInfo.h |
| headers/eventInfo.h |
| headers/openclBackend/openclBackendSynopsis.h |
| headers/openclBackend/createModuleWithSource.h |
| headers/openclBackend/createModuleWithBinary.h |
| headers/openclBackend/createModuleWithIL.h |
| headers/openclBackend/createModuleWithBuiltinKernels.h |
| headers/openclcInterop.h |

LISTINGS

1. Acknowledgements

Editors

- Maria Rovatsou, Codeplay
- Lee Howes, Qualcomm
- Ronan Keryell, Xilinx (current)

Contributors

- Eric Berdahl, Adobe
- Shivani Gupta, Adobe
- David Neto, Altera
- Brian Sumner, AMD
- Hal Finkel, Argonne National Laboratory
- Nevin Liber, Argonne National Laboratory
- Anastasia Stulova, ARM
- Balázs Keszthelyi, Broadcom
- Stuart Adams, Codeplay
- Gordon Brown, Codeplay
- Morris Hafner, Codeplay
- Alexander Johnston, Codeplay
- Marios Katsigiannis, Codeplay
- Paul Keir, Codeplay
- Steffen Larsen, Codeplay
- Victor Lomüller, Codeplay
- Tomas Matheson, Codeplay
- Duncan McBain, Codeplay

- Ralph Potter, Codeplay
- Ruyman Reyes, Codeplay
- Andrew Richards, Codeplay
- Maria Rovatsou, Codeplay
- Panagiotis Stratis, Codeplay
- Michael Wong, Codeplay
- Peter Žužek, Codeplay
- Matt Newport, EA
- Ruslan Arutyunyan, Intel
- Alexey Bader, Intel
- James Brodman, Intel
- Ilya Burylov, Intel
- Felipe de Azevedo Piovezan, Intel
- Allen Hux, Intel
- Michael Kinsner, Intel
- · Greg Lueck, Intel
- John Pennycook, Intel
- Roland Schulz, Intel

- Sergey Semenov, Intel
- Jason Sewall, Intel
- Kathleen Mattson, Miller & Mattson, LLC
- Dave Miller, Miller & Mattson, LLC
- Lee Howes, Qualcomm
- Chu-Cheow Lim, Qualcomm
- Jack Liu, Qualcomm
- Ruihao Zhang, Qualcomm
- Dave Airlie, Red Hat

- Aksel Alpay, Self
- Dániel Berényi, Self
- Máté Nagy-Egri, Stream HPC
- Tom Deakin, University of Bristol
- Paul Preney, University of Windsor
- Andrew Gozillon, Xilinx
- Ronan Keryell, Xilinx
- Lin-Ya Yu, Xilinx

2. Introduction

SYCL (pronounced "sickle") is a royalty-free, cross-platform abstraction C++ programming model for heterogeneous computing. SYCL builds on the underlying concepts, portability and efficiency of parallel API or standards like OpenCL while adding much of the ease of use and flexibility of single-source C++.

Developers using SYCL are able to write standard modern C++ code, with many of the techniques they are accustomed to, such as inheritance and templates. At the same time developers have access to the full range of capabilities of the underlying implementation (such as OpenCL) both through the features of the SYCL libraries and, where necessary, through interoperation with code written directly using the underneath implementation, via their APIs.

To reduce programming effort and increase the flexibility with which developers can write code, SYCL extends the concepts found in standards like OpenCL model in two ways beyond the general use of C++ features:

- Execution of parallel kernels on a heterogeneous device is made simultaneously convenient and flexible. Common parallel patterns are prioritised with simple syntax, which through a series C++ types allow the programmer to express additional requirements, such as synchronization, if needed.
- Data access in SYCL is separated from data storage. By relying on the C++-style resource acquisition is initialization (RAII) idiom to capture data dependencies between device code blocks, the runtime library can track data movement and provide correct behavior without the complexity of manually managing event dependencies between kernel instances and without the programming having to explicitly move data. This approach enables the data-parallel task-graphs that might be already part of the execution model to be built up easily and safely by SYCL programmers. To primarily assist with porting codes, SYCL provides an alternative way to manage data using Unified Shared Memory (USM).
- The hierarchical parallelism syntax offers a way of expressing the data-parallel similar to the OpenCL device or OpenMP target device execution model in an easy-to-understand modern C++ form. It more cleanly layers parallel loops and synchronization points to avoid fragmentation of code and to more efficiently map to CPU-style architectures.

SYCL retains the execution model, runtime feature set and device capabilities inspired by the OpenCL standard. This standard imposes some limitations on the full range of C++ features that SYCL is able to support. This ensures portability of device code across as wide a range of devices as possible. As a result, while the code can be written in standard C++ syntax with interoperability with standard C++ programs, the entire set of C++ features is not available in SYCL device code. In particular, SYCL device code, as defined by this specification, does not support virtual function calls, function pointers in general, exceptions, runtime type information or the full set of C++ libraries that may depend on these features or on features of a particular host compiler. Nevertheless, these basic restrictions can be relieved by some specific Khronos or vendor extensions.

SYCL implements a single-source multiple compiler-passes (SMCP) design which offers the power of source integration while allowing toolchains to remain flexible. The SMCP design supports embedding of code intended to be compiled for a device, for example a GPU, inline with host code. This embedding of code offers three primary benefits:

- **Simplicity** For novice programmers using frameworks like OpenCL, the separation of host and device source code in OpenCL can become complicated to deal with, particularly when similar kernel code is used for multiple different operations on different data types. A single compiler flow and integrated tool chain combined with libraries that perform a lot of simple tasks simplifies initial OpenCL programs to a minimum complexity. This reduces the learning curve for programmers new to heterogeneous programming and allows them to concentrate on parallelization techniques rather than syntax.
- **Reuse** C++'s type system allows for complex interactions between different code units and supports efficient abstract interface design and reuse of library code. For example, a *transform* or *map* operation applied to an array of data may allow specialization on both the operation applied to each element of the array and on the type of the data. The SMCP design of SYCL enables this interaction to bridge the host code/device code boundary such that the device code to be specialized on both of these factors directly from the host code.
- **Efficiency** Tight integration with the type system and reuse of library code enables a compiler to perform inlining of code and to produce efficient specialized device code based on decisions made in the host code without having to generate kernel source strings dynamically.

The use of C++ features such as generic programming, templated code, functional programming and inheritance on top of existing heterogeneous execution model opens a wide scope for innovation in software design for heterogeneous systems. Clean integration of device and host code within a single C++ type system enables the development of modern, templated generic and adaptable libraries that build simple, yet efficient, interfaces to offer more developers access to heterogeneous computing capabilities and devices. SYCL is intended to serve as a foundation for innovation in programming models for heterogeneous systems, that builds on open and widely implemented standard foundation like OpenCL or Vulkan.

SYCL is designed to be as close to standard C++ as possible. In practice, this means that as long as no dependence is created on SYCL's integration with the underlying implementation, a standard C++ compiler can compile the SYCL programs and they will run correctly on a host CPU. Any use of specialized low-level features can be masked using the C pre-processor in the same way that compiler-specific intrinsics may be hidden to ensure portability between different host compilers.

SYCL is designed to allow a compilation flow where the source file is passed through multiple different compilers, including a standard C++ host compiler of the developer's choice, and where the resulting application combines the results of these compilation passes. This is distinct from a single-source flow that might use language extensions that preclude the use of a standard host compiler. The SYCL standard does not preclude the use of a single compiler flow, but is designed to not require it.

The advantages of this design are two-fold. First, it offers better integration with existing tool chains. An application that already builds using a chosen compiler can continue to do so when SYCL code is added. Using the SYCL tools on a source file within a project will both compile for a device and let the same source file be compiled using the same host compiler that the rest of the project is compiled with. Linking and library relationships are unaffected. This design simplifies porting of pre-existing applications to SYCL. Second, the design allows the optimal compiler to be chosen for each device where different vendors may provide optimized tool-chains.

To summarize, SYCL enables computational kernels to be written inside C++ source files as normal C++ code, leading to the concept of "single-source" programming. This means that software developers can develop and use generic algorithms and data structures using standard C++ template techniques, while still supporting multiplatform, multi-device heterogeneous execution. Access to the low level APIs of an underlying implementation (such as OpenCL) is also supported. The specification has been designed to enable implementation across as wide a variety of platforms as possible as well as ease of integration with other platform-specific technologies, thereby letting both users and implementers build on top of SYCL as an open platform for system-wide heterogeneous processing innovation.

3. SYCL architecture

This chapter describes the structure of a SYCL application, and how the SYCL generic programming model lays out on top of a number of SYCL backends.

3.1 Overview

SYCL is an open industry standard for programming a heterogeneous system. The design of SYCL allows standard C++ source code to be written such that it can run on either an heterogeneous device or on the host.

The terminology used for SYCL inherits historically from OpenCL with some SYCL-specific additions. However SYCL is a generic C++ programming model that can be layed out on top of other heterogeneous APIs apart from OpenCL. SYCL implementations can provide SYCL backends for various heterogeneous APIs, implementing the SYCL general specification on top of them. We refer to this heterogeneous API as the SYCL backend API. The SYCL general specification defines the behavior that all SYCL implementations must expose to SYCL users for a SYCL application to behave as expected.

A function object that can execute on a device exposed by a SYCL backend API is called a SYCL kernel function.

To ensure maximum interoperability with different SYCL backend APIs, software developers can access the SYCL backend API alongside the SYCL general API whenever they include the SYCL backend interoperability headers. However, interoperability is a SYCL backend-specific feature. An application that uses interoperability does not conform to the SYCL general application model, since it is not portable across backends.

The target users of SYCL are C++ programmers who want all the performance and portability features of OpenCL, but with the flexibility to use higher-level C++ abstractions across the host/device code boundary. Developers can use most of the abstraction features of C++, such as templates, classes and operator overloading.

However, some C++ language features are not permitted inside kernels, due to the limitations imposed by the capabilities of the underlying heterogeneous platforms. These features include virtual functions, virtual inheritance, throwing/catching exceptions, and run-time type-information. These features are available outside kernels as normal. Within these constraints, developers can use abstractions defined by SYCL, or they can develop their own on top. These capabilities make SYCL ideal for library developers, middleware providers and application developers who want to separate low-level highly-tuned algorithms or data structures that work on heterogeneous systems from higher-level software development. Software developers can produce templated algorithms that are easily usable by developers in other fields.

3.2 Anatomy of a SYCL application

Below is an example of a typical SYCL application which schedules a job to run in parallel on any heterogeneous device available.

```
#include <SYCL/sycl.hpp>
 1
 2
    #include <iostream>
 3
 4
    int main() {
 5
      using namespace sycl;
 6
 7
      int data[1024]; // Allocate data to be worked on
 8
 9
      // By wrapping all the SYCL work in a {} block, we ensure
10
      // all SYCL tasks must complete before exiting the block,
11
      // because the destructor of resultBuf will wait.
12
      {
13
        // Create a queue to enqueue work to
14
        queue myQueue;
15
16
        // Wrap our data variable in a buffer
17
        buffer<int, 1> resultBuf { data, range<1> { 1024 } };
18
19
        // Create a command_group to issue commands to the queue
20
        myQueue.submit([&](handler& cgh) {
21
          // request access to the buffer
22
          accessor writeResult { resultBuf, cgh, write_only, noinit };
23
24
          // Enqueue a parallel_for task
25
          cgh.parallel_for(1024, [=](auto idx) {
26
            writeResult[idx] = idx;
27
          }); // End of the kernel function
28
               // End of our commands for this queue
        });
29
      }
               // End of scope, so we wait for work producing resultBuf to complete
30
31
      // Print result
32
      for (int i = 0; i < 1024; i++)
33
        std::cout << "data[" << i << "] = " << data[i] << std::endl;</pre>
34
35
      return 0;
36 }
```

At line 1, we "#include" the SYCL header files, which provide all of the SYCL features that will be used.

A SYCL application runs on a SYCL Platform (see Section 3.4). The application is structured in three scopes which specify the different sections; application scope, command group scope and kernel scope. The kernel scope specifies a single kernel function that will be, or has been, compiled by a device compiler and executed on a device. In this example kernel scope is defined by lines 25 to 27. The command group scope specifies a unit of work which is comprised of a SYCL kernel function and accessors. In this example command group scope is defined by lines 20 to 28. The application scope specifies all other code outside of a command group scope. These three scopes are used to control the application flow and the construction and lifetimes of the various objects used within SYCL, as explained in Section 3.8.12.

A SYCL kernel function is the scoped block of code that will be compiled using a device compiler. This code may be defined by the body of a lambda function or by the operator() function of a function object. Each instance of the SYCL kernel function will be executed as a single, though not necessarily entirely independent, flow of execution and has to adhere to restrictions on what operations may be allowed to enable device compilers to safely compile it to a range of underlying devices.

The parallel_for function can be templated with a class. This class is used to manually name the kernel when desired, such as to avoid a compiler-generated name when debugging a kernel defined through a lambda, to provide a known name with which to apply build options to a kernel, or to ensure compatibility with multiple compiler-pass implementations.

The parallel_for member function creates an instance of a kernel object. The kernel object is the entity that will be enqueued within a command group. In the case of parallel_for the SYCL kernel function will be executed over the given range from 0 to 1023. The different member functions to execute kernels can be found in Section 4.10.7.

A command group scope is the syntactic scope wrapped by the construction of a command group function object as seen on line 20. The command group function object may invoke only a single SYCL kernel function, and it takes a parameter of type command group handler, which is constructed by the runtime.

All the requirements for a kernel to execute are defined in this command group scope, as described in Section 3.6.1. In this case the constructor used for myQueue on line 14 is the default constructor, which allows the queue to select the best underlying device to execute on, leaving the decision up to the runtime.

In SYCL, data that is required within a SYCL kernel function must be contained within a buffer or image, as described in Section 3.7. We construct a buffer on line 17. Access to the buffer is controlled via an accessor which is constructed on line 22 through the get_access member function of the buffer. The buffer is used to keep track of access to the data and the accessor is used to request access to the data on a queue, as well as to track the dependencies between SYCL kernel function. In this example the accessor is used to write to the data buffer on line 26. All buffers must be constructed in the application-scope, whereas all accessors must be constructed in the command group scope.

3.3 Normative references

The documents in the following list are referred to within this SYCL specification, and their content is a requirement for this document.

- 1. C++17: ISO/IEC 14882:2017 Clauses 1-19 [2], referred to in this specification as the C++ core language.
- 2. C++2a: Working Draft, Standard for Programming Language C++ [3], referred to in this specification as the next C++ specification.

3.4 The SYCL platform model

The SYCL platform model is based on the OpenCL platform model. The model consists of a host connected to one or more heterogeneous devices, called devices.

A SYCL context is constructed, either directly by the user or implicitly when creating a queue, to hold all the runtime information required by the SYCL runtime and the SYCL backend to operate on a device, or group of devices. When a group of devices can be grouped together on the same context, they have some visibility of each other's memory objects. The SYCL runtime can assume that memory is visible across all devices in the same

context. Not all devices exposed from the same platform can be grouped together in the same context.

A SYCL application executes on the host as a standard C++ program. Devices are exposed through different SYCL backends to the SYCL application. The SYCL application submits command group function objects to queues. Each queue enables execution on a given device.

The SYCL runtime then extracts operations from the command group function object, e.g. an explicit copy operation or a SYCL kernel function. When the operation is a SYCL kernel function, the SYCL runtime uses a SYCL backend-specific mechanism to extract the device binary from the SYCL application and pass it to the heterogeneous API for execution on the device.

A SYCL device is divided into one or more compute units (CUs) which are each divided into one or more processing elements (PEs). Computations on a device occur within the processing elements. How computation is mapped to PEs is SYCL backend and device specific. Two devices exposed via two different backends can map computations differently to the same device.

When a SYCL application contains SYCL kernel function objects, the SYCL implementation must provide an offline compilation mechanism that enables the integration of the device binaries into the SYCL application. The output of the offline compiler can be an intermediate representation, such as SPIR-V, that will be finalized during execution or a final device ISA.

A device may expose special purpose functionality as a *built-in* function. The SYCL API exposes functions to query and dispatch said *built-in* functions. Some SYCL backends and devices may not support programmable kernels, and only support *built-in* functions.

3.5 The SYCL backend model

SYCL is a generic programming model for the C++ language that can target multiple heterogeneous APIs, such as OpenCL.

SYCL implementations enable these target APIs by implementing SYCL backends. For a SYCL implementation to be conformant on said SYCL backend, it must execute the SYCL generic programming model on the backend.

All SYCL implementations must provide a host SYCL backend, which implements the SYCL API and executes SYCL kernel functions on the host platform. Many implementations will provide additional SYCL backends that execute SYCL kernel functions on an accelerator.

The present document covers the SYCL generic interface available to all SYCL backends. How the SYCL generic interface maps to a particular SYCL backend is defined either by a separate SYCL backend specification document, provided by the Khronos SYCL group, or by the SYCL implementation documentation. Whenever there is a SYCL backend specification document, this takes precedence over SYCL implementation documentation.

When a SYCL user builds their SYCL application, she decides which of the SYCL backends will be used to build the SYCL application. This is called the set of *active backends*. Implementations must ensure that the active backends selected by the user can be used simultaneously by the SYCL implementation at runtime. If two backends are available at compile time but will produce an invalid SYCL application at runtime, the SYCL implementation must emit a compilation error.

A SYCL application built with a number of active backends does not necessarily guarantee that said backends can be executed at runtime. The subset of active backends available at runtime is called *available backends*. A backend is said to be *available* if the host platform where the SYCL application is executed exposes support for the heterogeneous API required for the SYCL backend.

It is implementation dependent whether certain backends require third-party libraries to be available in the system. Building with only the host as an active backend guarantees the binary will be executed on any platform without requiring third-party libraries. Failure to have all dependencies required for all active backends at runtime will cause the SYCL application to not run.

Once the application is running, users can query what SYCL platforms are available. SYCL implementations will expose the devices provided by each backend grouped into platforms. A backend must expose at least one platform.

Under the SYCL backend model, SYCL objects can contain one or multiple references to a certain SYCL backend native type. Not all SYCL objects will map directly to a SYCL backend native type. The mapping of SYCL objects to SYCL backend native types is defined by the SYCL backend specification document when available, or by the SYCL implementation otherwise.

To guarantee that multiple SYCL backend objects can interoperate with each other, SYCL memory objects are not bound to a particular SYCL backend. SYCL memory objects can be accessed from any device exposed by an *available* backend. SYCL Implementations can potentially map SYCL memory objects to multiple native types in different SYCL backends.

Since SYCL memory objects are independent of any particular SYCL backend, SYCL command groups can request access to memory objects allocated by any SYCL backend, and execute it on the backend associated with the queue. This requires the SYCL implementation to be able to transfer memory objects across SYCL backends

When a SYCL application runs on any number of SYCL backends without relying on any SYCL backend-specific behaviour or interoperability, it is said to be a SYCL general application, and it is expected to run in any SYCL-conformant implementation that supports the required features for the application.

3.5.1 Platform mixed version support

The SYCL generic programming model exposes a number of platforms, each of them exposing a number of devices. Each platform is bound to a certain SYCL backend. SYCL devices associated with said platform are associated with that SYCL backend.

Although the APIs in the SYCL generic programming model are defined according to this specification and their version is indicated by the macro SYCL_LANGUAGE_VERSION, this does not apply to APIs exposed by the SYCL backends. Each SYCL backend provides its own document that defines its APIs, and that document tells how to query for the device and platform versions.

3.6 SYCL execution model

As Section 3.2 describes, a SYCL application is comprised of three scopes: application scope, command group scope, and kernel scope. Code in the application scope and command group scope runs on the host and is governed by the *SYCL application execution model*. Code in the kernel scope runs on a device and is governed by the *SYCL kernel execution model*.

A SYCL application can be executed on the host directly without a physical accelerator present when using the SYCL host backend.

3.6.1 SYCL application execution model

The SYCL application defines the execution order of the kernels by grouping each kernel with its requirements into a command group function object. Command group function objects are submitted to execution via a queue object, which defines the device where the kernel will run. The same command group object can be submitted to different queues. When a command group is submitted to a SYCL queue, the requirements of the kernel execution are captured. The kernels are executed as soon as their requirements have been satisfied.

3.6.1.1 SYCL backend resources managed by the SYCL application

The SYCL runtime integrated with the SYCL application will manage the resources required by the SYCL backend API to manage the heterogeneous devices it is providing access to. This includes, but is not limited to, resource handlers, memory pools, dispatch queues and other temporary handler objects.

The SYCL programming interface represents the lifetime of the resources managed by the SYCL application using RAII rules. Construction of a SYCL object will typically entail the creation of multiple SYCL backend objects, which will be properly released on destruction of said SYCL object. The overall rules for construction and destruction are detailed in the SYCL Programming Interface chapter 4. Those SYCL backends with a SYCL backend document will detail how the resource management from SYCL objects map down to the SYCL backend objects.

In SYCL, the minimum required object for submitting work to devices is the queue, which contains references to a platform, device and a context internally.

The resources managed by SYCL are:

- 1. Platforms: all features of SYCL backend APIs are implemented by platforms. A platform can be viewed as a given vendor's runtime and the devices accessible through it. Some devices will only be accessible to one vendor's runtime and hence multiple platforms may be present. SYCL manages the different platforms for the user.
- 2. Contexts: any SYCL backend resource that is acquired by the user is attached to a context. A context contains a collection of devices that the host can use and manages memory objects that can be shared between the devices. Devices belonging to the same context must be able to access each other's global memory using some implementation-specific mechanism. A given context can only wrap devices owned by a single platform.
- 3. Devices: platforms provide one or more devices for executing SYCL kernels. In SYCL, a device is accessible through a sycl::device object.
- 4. Kernels: the SYCL functions that run on SYCL devices are defined as C++ function objects (a named function object type or a lambda function).

Note that some SYCL backends may expose non-programmable functionality as pre-defined kernels.

- 5. Modules: The SYCL backend API will expose the device binaries in some form of a program-object file or module. This is handled by SYCL using the module objects.
- 6. Queues: SYCL kernels execute in command queues. The user must create a queue, which references an associated context, platform and device. The context, platform and device may be chosen automatically, or specified by the user. SYCL queues execute kernels on a particular device of a particular context, but can have dependencies from any device on any available SYCL backend.

The SYCL implementation guarantees the correct initialization and destruction of any resource handled by the underlying SYCL backend API, except for those the user has obtained manually via the SYCL interoperability API.

3.6.1.2 SYCL command groups and execution order

By default, SYCL queues execute kernel functions in an out-of-order fashion based on dependency information. Developers only need to specify what data is required to execute a particular kernel. The SYCL runtime will guarantee that kernels are executed in an order that guarantees correctness. By specifying access modes and types of memory, a directed acyclic dependency graph (DAG) of kernels is built at runtime. This is achieved via the usage of command group objects. A SYCL command group object defines a set of requisites (*R*) and a kernel function (*k*). A command group is *submitted* to a queue when using the sycl::queue::submit member function.

A **requisite** (r_i) is a requirement that must be fulfilled for a kernel-function (k) to be executed on a particular device. For example, a requirement may be that certain data is available on a device, or that another command group has finished execution. An implementation may evaluate the requirements of a command group at any point after it has been submitted. The *processing of a command group* is the process by which a SYCL runtime evaluates all the requirements in a given *R*. The SYCL runtime will execute *k* only when all r_i are satisfied (i.e., when all requirements are satisfied). To simplify the notation, in the specification we refer to the set of requirements of a command group named *foo* as $CG_{foo} = r_1, \ldots, r_n$.

The *evaluation of a requisite* (Satisfied(r_i)) returns the status of the requisite, which can be *True* or *False*. A *satisfied* requisite implies the requirement is met. Satisfied(r_i) never alters the requisite, only observes the current status. The implementation may not block to check the requisite, and the same check can be performed multiple times.

An **action** (a_i) is a collection of implementation-defined operations that must be performed in order to satisfy a requisite. The set of actions for a given command group A is permitted to be empty if no operation is required to satisfy the requirement. The notation a_i represents the action required to satisfy r_i . Actions of different requisites can be satisfied in any order w.r.t each other without side effects (i.e., given two requirements r_j and r_k , $(r_j, r_k) \equiv (r_k, r_j)$). The intersection of two actions is not necessarily empty. Actions can include (but are not limited to): memory copy operations, mapping operations, host side synchronization, or implementation-specific behavior.

Finally, *Performing an action* (Perform(a_i)) executes the action operations required to satisfy the requisite r_j . Note that, after Perform(a_i), the evaluation Satisfied(r_j) will return *True* until the kernel is executed. After the kernel execution, it is not defined whether a different command group with the same requirements needs to perform the action again, where actions of different requisites inside the same command group object can be satisfied in any order w.r.t each other without side effects: Given two requirements r_j and r_k , Perform(a_j) followed by Perform(a_k) is equivalent to Perform(a_k) followed by Perform(a_j).

The requirements of different command groups submitted to the same or different queues are evaluated in the relative order of submission. command group objects whose intersection of requirement sets is not empty are said to depend on each other. They are executed in order of submission to the queue. If command groups are submitted to different queues or by multiple threads, the order of execution is determined by the SYCL runtime. Note that independent command group objects can be submitted simultaneously without affecting dependencies.

Figure 3.1 illustrates the execution order of three command group objects (CG_a, CG_b, CG_c) with certain requirements submitted to the same queue. Both CG_a and CG_b only have one requirement, r_1 and r_2 respectively. CG_c requires both r_1 and r_2 . This enables the SYCL runtime to potentially execute CG_a and CG_b simultaneously, whereas CG_c cannot be executed until both CG_a and CG_b have been completed. The SYCL runtime evaluates the **requisites** and performs the **actions** required (if any) for the CG_a and CG_b . When evaluating the **requisites** of CG_c , they will be satisfied once the CG_a and CG_b have finished.

SYCL Application Enqueue Order

SYCL Kernel Execution Order



Figure 3.1: Execution order of three command groups submitted to the same queue.

Figure 3.2 uses three separate SYCL queue objects to submit the same command group objects as before. Regardless of using three different queues, the execution order of the different command group objects is the same. When different threads enqueue to different queues, the execution order of the command group will be the order in which the submit member function is executed. In this case, since the different command group objects execute on different devices, the **actions** required to satisfy the **requirements** may be different (e.g, the SYCL runtime may need to copy data to a different device in a separate context).

SYCL Application Enqueue OrderSYCL Kernel Execution Order $sycl::queue syclQueue1; sycl::queue syclQueue3; syclQueue2; sycl::queue syclQueue3; syclQueue2.submit (<math>CG_a(r_1)$);
 $syclQueue3.submit (<math>CG_c(r_1, r_2)$); $CG_a(r_1)$

Figure 3.2: Execution order of three command groups submitted to the different queues.

3.6.2 SYCL kernel execution model

When a kernel is submitted for execution, an index space is defined. An instance of the kernel body executes for each point in this index space. This kernel instance is called a work-item and is identified by its point in the index space, which provides a global id for the work-item. Each work-item executes the same code but the specific execution pathway through the code and the data operated upon can vary by using the work-item global id to specialize the computation.

Work-items are organized into work-groups. The work-groups provide a more coarse-grained decomposition of the index space. Each work-group is assigned a unique work-group id with the same dimensionality as the index space used for the work-items. Work-items are each assigned a local id, unique within the work-group, so that a single work-item can be uniquely identified by its global id or by a combination of its local id and work-group id. The work-items in a given work-group execute concurrently on the processing elements of a single compute unit.

The index space supported in SYCL is called an nd-range. An ND-range is an N-dimensional index space, where N is one, two or three. In SYCL, the ND-range is represented via the nd_range<N> class. An nd_range<N> is
made up of a global range and a local range, each represented via values of type range<N> and a global offset, represented via a value of type id<N>. The types range<N> and id<N> are each N-element arrays of integers. The iteration space defined via an nd_range<N> is an N-dimensional index space starting at the ND-range's global offset whose size is its global range, split into work-groups of the size of its local range.

Each work-item in the ND-range is identified by a value of type nd_item<N>. The type nd_item<N> encapsulates a global id, local id and work-group id, all of type id<N>, the iteration space offset also of type id<N>, as well as global and local ranges and synchronization operations necessary to make work-groups useful. Work-groups are assigned ids using a similar approach to that used for work-item global ids. Work-items are assigned to a work-group and given a local id with components in the range from zero to the size of the work-group in that dimension minus one. Hence, the combination of a work-group id and the local id within a work-group uniquely defines a work-item.

SYCL allows a simplified execution model in which the work-group size is left unspecified. A kernel invoked over a range<N>, instead of an nd_range<N> is executed within an iteration space of unspecified work-group size. In this case, less information is available to each work-item through the simpler item<N> class.

SYCL allows SYCL backends to expose fixed functionality as non-programmable kernels. The behavior of these functions are SYCL backend specific, and do not necessarily follow the SYCL kernel execution model.

3.7 Memory model

Since SYCL is a single-source programming model, the memory model affects both the application and the device kernel parts of a program. On the SYCL application, the SYCL runtime will make sure data is available for execution of the kernels. On the SYCL device kernel, the SYCL backend rules describing how the memory behaves on a specific device are mapped to SYCL C++ constructs. Thus it is possible to program kernels efficiently in pure C++.

3.7.1 SYCL application memory model

The application running on the host uses SYCL buffer objects using instances of the sycl::buffer class to allocate memory in the global address space, or can allocate specialized image memory using the sycl::unsampled_image and sycl::sampled_image classes.

In the SYCL application, memory objects are bound to all devices in which they are used, regardless of the SYCL context where they reside. SYCL memory objects (namely, buffer and image objects) can encapsulate multiple underlying SYCL backend memory objects together with multiple host memory allocations to enable the same object to be shared between devices in different contexts, platforms or backends.

The order of execution of command group objects ensures a sequentially consistent access to the memory from the different devices to the memory objects.

To access a memory object, the user must create an accessor object which parameterizes the type of access to the memory object that a kernel or the host requires. The accessor object defines a requirement to access a memory object, and this requirement is defined by construction of an accessor, regardless of whether there are any uses in a kernel or by the host. The cl::sycl::accessor object specifies whether the access is via global memory, constant memory or image samplers and their associated access functions. The accessor also specifies whether the access is read-only (RO), write-only (WO) or read-write (RW). An optional *discard* flag can be added to an accessor to tell the system to discard any previous contents of the data the accessor refers to, e.g. discard write-only (DW). For simplicity, when a **requisite** represents an accessor object in a certain access mode, we represent it as MemoryObject_{AccessMode}. For example, an accessor that accesses memory object **buf1** in **RW** mode is represented

as $buf 1_{RW}$. A command group object that uses such an accessor is represented as $CG(buf 1_{RW})$. The **action** required to satisfy a requisite and the location of the latest copy of a memory object will vary depending on the implementation.

Figure 3.3 illustrates an example where command group objects are enqueued to two separate SYCL queues executing in devices in different contexts. The **requisites** for the command group execution are the same, but the **actions** to satisfy them are different. For example, if the data is on the host before execution, $A(b1_{RW})$ and $A(b2_{RW})$ can potentially be implemented as copy operations from the host memory to context1 or context2 respectively. After CG_a and CG_b are executed, $A'(b1_{RW})$ will likely be an empty operation, since the result of the kernel can stay on the device. On the other hand, the results of CG_b are now on a different context than CG_c is executing, therefore $A'(b2_{RW})$ will need to copy data across two separate OpenCL contexts using an implementation specific mechanism.

SYCL Application Enqueue Order

SYCL Kernel Execution Order



Possible implementation by a SYCL Runtime



Figure 3.3: Actions performed when three command groups are submitted to two distinct queues, and potential implementation in an OpenCL SYCL backend by a SYCL runtime. Note that in this example, each SYCL buffer (b1, b2) is implemented as separate cl_mem objects per context.

Note that the order of the definition of the accessors within the command group is irrelevant to the requirements they define. All accessors always apply to the entire command group object where they are defined.

When multiple accessors in the same command group define different requisites to the same memory object these requisites must be resolved.

Firstly, any requisites with different access modes but the same access target are resolved into a single requisite with the union of the different access modes according to Table 3.1. The atomic access mode acts as if it was read-

| One access mode | Other access mode | Combined requirement |
|--------------------------|--------------------------|--------------------------|
| read (RO) | write (WO) | read-write (RW) |
| read (RO) | read-write (RW) | read-write (RW) |
| write (WO) | read-write (RW) | read-write (RW) |
| discard-write (DW) | discard-read-write (DRW) | discard-read-write (DRW) |
| discard-write (DW) | write (WO) | write (WO) |
| discard-write (DW) | read (RO) | read-write (RW) |
| discard-write (DW) | read-write (RW) | read-write (RW) |
| discard-read-write (DRW) | write (WO) | read-write (RW) |
| discard-read-write (DRW) | read (RO) | read-write (RW) |
| discard-read-write (DRW) | read-write (RW) | read-write (RW) |

write (RW) when determining the combined requirement. The rules in Table 3.1 are commutative and associative.

Table 3.1: Combined requirement from two different accessor access modes within the same command group. The rules are commutative and associative.

The result of this should be that there should not be any requisites with the same access target.

Secondly, the remaining requisites must adhere to the following rule. Only one of the requisites may have write access (W or RW), otherwise the SYCL runtime must throw an exception. All requisites create a requirement for the data they represent to be made available in the specified access target, however only the requisite with write access determines the side effects of the command group, i.e. only the data which that requisite represents will be updated.

For example:

- $CG(b1_{RW}^G, b1_R^H)$ is permitted.
- $CG(b1_{RW}^G, b1_{RW}^H)$ is **not** be permitted.
- $CG(b1_W^G, b1_{RW}^C)$ is **not** be permitted.

Where *G*, *C* and *H* correspond to the global_buffer, constant_buffer and host_buffer respectively.

A buffer created from a range of an existing buffer is called a *sub-buffer*. A buffer may be overlaid with any number of sub-buffers. Accessors can be created to operate on these *sub-buffers*. Refer to 4.7.2 for details on *sub-buffer* creation and restrictions. A requirement to access a sub-buffer is represented by specifying its range, e.g. $CG(b_{1_{RW[0,5)}})$ represents the requirement of accessing the range [0, 5) buffer b1 in read write mode.

If two accessors are constructed to access the same buffer, but both are to non-overlapping sub-buffers of the buffer, then the two accessors are said to not *overlap*, otherwise the accessors do overlap. Overlapping is the test that is used to determine the scheduling order of command groups. Command-groups with non-overlapping requirements may execute concurrently.

It is permissible for command groups that only read data to not copy that data back to the host or other devices after reading and for the runtime to maintain multiple read-only copies of the data on multiple devices.

A special case of requirement is the one defined by a **host accessor**. Host accessors are represented with $H(\text{MemoryObject}_{\text{accessMode}})$, e.g, $H(b1_{RW})$ represents a host accessor to b1 in read-write mode. Host accessors are a special type of accessor constructed from a memory object outside a command group, and require that the data associated with the given memory object is available on the host in the given pointer. This causes the runtime to block on construction of this object until the requirement has been satisfied. **Host accessor** objects are effected.

SYCL Application Enqueue Order

SYCL Kernel Execution Order



Figure 3.4: Requirements on overlapping vs non-overlapping sub-buffer.

tively barriers on all accesses to a certain memory object. Figure 3.5 shows an example of multiple command groups enqueued to the same queue. Once the host accessor $H(b1_{RW})$ is reached, the execution cannot proceed until CG_a is finished. However, CG_b does not have any requirements on b1, therefore, it can execute concurrently with the barrier. Finally, CG_c will be enqueued after $H(b1_{RW})$ is finished, but still has to wait for CG_b to conclude for all its requirements to be satisfied. See 3.8.8 for details on synchronization rules.

SYCL Application Enqueue Order

SYCL Kernel Execution Order

sycl::queue q1; q1.submit ($CG_a(b1_{RW})$); q1.submit ($CG_b(b2_{RW})$); $H(b1_{RW})$; q1.submit ($CG_c(b1_{RW}, b2_{RW})$);



Figure 3.5: Execution of command groups when using host accessors.

3.7.2 SYCL device memory model

The memory model for SYCL devices is based on the OpenCL 1.2 memory model. Work-items executing in a kernel have access to four distinct address spaces (memory regions) and a virtual address space overlapping some concrete address spaces:

- Global memory is accessible to all work-items in all work-groups. Work-items can read from or write to
 any element of a global memory object. Reads and writes to global memory may be cached depending on
 the capabilities of the device. Global memory is persistent across kernel invocations, however there is no
 guarantee that two concurrently executing kernels can simultaneously write to the same memory object and
 expect correct results.
- Constant memory is a region of global memory that remains constant during the execution of a kernel. The host allocates and initializes memory objects placed into constant memory.

- Local memory is accessible to all work-items in a single work-group and inaccessible to work-items in other work-groups. This memory region can be used to allocate variables that are shared by all work-items in a work-group. Work-group-level visibility allows local memory to be implemented as dedicated regions of the device memory where this is appropriate.
- Private memory is a region of memory private to a work-item. Variables defined in one work-item's private memory are not visible to another work-item.
- Generic memory is a virtual address space which overlaps the global, local and private address spaces.

3.7.2.1 Access to memory

Accessors in the device kernels provide access to the memory objects, acting as pointers to the corresponding address space.

It is not possible to pass a pointer into host memory directly as a kernel parameter because the devices may be unable to support the same address space as the host.

To allocate local memory within a kernel, the user can either pass a sycl::local_accessor object to the kernel as a parameter, or can define a variable in work-group scope inside sycl::parallel_for_work_group.

Any variable defined inside a sycl::parallel_for scope or sycl::parallel_for_work_item scope will be allocated in private memory. Any variable defined inside a sycl::parallel_for_work_group scope will be allocated in local memory.

Users can create accessors that reference sub-buffers as well as entire buffers.

Within kernels, the underlying C++ pointer types can be obtained from an accessor. The pointer types will contain a compile-time deduced address space. So, for example, if a C++ pointer is obtained from an accessor to global memory, the C++ pointer type will have a global address space attribute attached to it. The address space attribute will be compile-time propagated to other pointer values when one pointer is initialized to another pointer value using a defined algorithm.

When developers need to explicitly state the address space of a pointer value, one of the explicit pointer classes can be used. There is a different explicit pointer class for each address space: sycl::raw_local_ptr, sycl::raw_global_ptr, sycl::raw_private_ptr, sycl::raw_constant_ptr, sycl::raw_generic_ptr , sycl::decorated_local_ptr, sycl::decorated_global_ptr, sycl::decorated_private_ptr, sycl:: decorated_constant_ptr, or sycl::decorated_generic_ptr. The classes with the decorated prefix expose pointers that use an implementation defined address space decoration, while the classes with the raw prefix do not. Accessors with an access target global_buffer, constant_buffer, or local, can be converted into explicit pointer classes (multi_ptr). Explicit pointer class values cannot be passed as parameters to kernels or stored in global memory.

For templates that need to adapt to different address spaces, a sycl::multi_ptr class is defined which is templated via a compile-time constant enumerator value to specify the address space.

3.7.2.2 Memory consistency inside SYCL kernels

The SYCL memory consistency model is based upon the memory consistency model of the C++ core language. Where SYCL offers extensions to classes and functions that may affect memory consistency (e.g. sycl::atomic_ref), the default behavior when these extensions are not used always matches the behavior of standard C++.

A SYCL implementation must guarantee that the same memory consistency model is used across host and device code. Every device compiler must support the memory model defined by the minimum version of C++ described in Section 3.8.1; SYCL implementations supporting additional versions of C++ must also support the corresponding memory models.

The full C++ memory model is not guaranteed to be supported by every device, nor across all combinations of devices within a context. The memory orderings supported by a specific device and context can be queried using functionalities of the sycl::device and sycl::context classes, respectively.

Within a SYCL kernel, all memory has load/store consistency within a work-item. Ensuring memory consistency across different work-items requires careful usage of group barrier operations, mem-fence operations and atomic operations. On any SYCL device, local and global memory may be made consistent across work-items in a single group through use of a group barrier operation. On SYCL devices supporting acquire-release or sequentially consistent memory orderings, all memory visible to a set of work-items may be made consistent (with a single *happens-before* relation) across the work-items in that set through the use of mem-fence and atomic operations.

The set of work-items and devices to which the memory ordering constraints of a given atomic operation apply is controlled by a memory scope constraint, which can take one of the following values:

- memory_scope::work_item The ordering constraint applies only to the calling work-item. This is only useful for image operations, since all other operations within a work-item are guaranteed to execute in program order.
- memory_scope::sub_group The ordering constraint applies only to work-items in the same sub-group as the calling work-item.
- memory_scope::work_group The ordering constraint applies only to work-items in the same work-group as the calling work-item. This is the broadest scope that can be applied to atomic operations in work-group local memory. Using any broader scope for atomic operations in work-group local memory is treated as though memory_scope::work_group was specified.
- memory_scope::device The ordering constraint applies only to work-items executing on the same device as the calling work-item.
- memory_scope::system The ordering constraint applies to any work-item or host thread in the system that is currently permitted to access the memory allocation containing the referenced object, as defined by the capabilities of buffers and USM.

[Note for this provisional version: The addition of memory scopes to the C++ memory model modifies the definition of some concepts from the C++ core language. For example: data races, the synchronizes-with relationship and sequential consistency must be defined in a way that accounts for atomic operations with differing (but compatible) scopes, in a manner similar to the OpenCL 2.0 specification [4]. Modified definitions of these concepts will be included in the final version of this specification. — end note]

These memory consistency guarantees are independent of any forward progress guarantees. Communication and synchronization between work-items in different work-groups is unsafe in general, but is supported on devices where all of the following conditions are true:

- Acquire-release or sequentially consistent memory ordering is supported at device scope
- · Work-items in different work-groups make independent forward progress

3.7.2.3 Atomic operations

Atomic operations can be performed on memory in buffers and USM. The range of atomic operations available on a specific device is limited by the atomic capabilities of that device. The sycl::atomic_ref class must be used to provide safe atomic access to the buffer or USM allocation from device code.

3.8 The SYCL programming model

A SYCL program is written in standard C++. Host code and device code is written in the same C++ source file, enabling instantiation of templated kernels from host code and also enabling kernel source code to be shared between host and device. The device kernels are encapsulated C++ callable types (a function object with operator () or a lambda function), which have been designated to be compiled as SYCL kernels.

SYCL programs target heterogeneous systems. The kernels may be compiled and optimized for multiple different processor architectures with very different binary representations.

3.8.1 Minimum version of C++

The C++ features used in SYCL are based on a specific version of C++. Implementations of SYCL must support this minimum C++ version, which defines the C++ constructs that can consequently be used by SYCL feature definitions (for example, lambdas).

The minimum C++ version of this SYCL specification is determined by the normative C++ core language defined in Section 3.3. All implementations of this specification must support at least this core language, and features within this specification are defined using features of the core language. Note that not all core language constructs are supported within SYCL kernel functions or code invoked by a SYCL kernel function, as detailed by Section 5.4.

Implementations may support newer C++ versions than the minimum required by SYCL. Code written using newer features than the SYCL requirement, though, may not be portable to other implementations that don't support the same C++ version.

3.8.2 Alignment with future versions of C++

Some features of SYCL are aligned with the next C++ specification, as defined in Section 3.3.

The following features are pre-adopted by SYCL 2020 and made available in the sycl:: namespace: std::span, std::bit_cast. The implementations of pre-adopted features are compliant with the next C++ specification, and are expected to forward directly to standard C++ features in a future version of SYCL.

The following features of SYCL 2020 use syntax based on the next C++ specification: sycl::atomic_ref. These features behave as described in the next C++ specification, barring modifications to ensure compatibility with other SYCL 2020 features and heterogeneous programming. Any such modifications are documented in the corresponding sections of this specification.

3.8.3 Basic data parallel kernels

Data-parallel kernels that execute as multiple work-items and where no local synchronization is required are enqueued with the sycl::parallel_for function parameterized by a sycl::range parameter. These kernels will execute the kernel function body once for each work-item in the range. The range passed to sycl::parallel_for

represents the global size of a SYCL kernel and will be divided into work-groups whose size is chosen by the SYCL runtime. Barrier synchronization is not valid within these work-groups.

Variables with reduction semantics can be added to basic data parallel kernels using the features described in Section 4.10.2.

3.8.4 Work-group data parallel kernels

Data parallel kernels can also execute in a mode where the set of work-items is divided into work-groups of user-defined dimensions. The user specifies the global range and local work-group size as parameters to the sycl ::parallel_for function with a sycl::nd_range parameter. In this mode of execution, kernels execute over the nd_range in work-groups of the specified size. It is possible to share data among work-items within the same work-group in local or global memory and to synchronize between work-items in the same work-group by calling the group_barrier function. All work-groups in a given parallel_for will be the same size, and the global size defined in the nd_range must be a multiple of the work-group size in each dimension.

Work-groups may be further subdivided into sub-groups. The size and number of sub-groups is implementationdefined and may differ for each kernel, and different devices may make different guarantees with respect to how sub-groups within a work-group are scheduled. The maximum number of work-items in any sub-group in a kernel is based on a combination of the kernel and its dispatch dimensions. The size of any sub-group in the dispatch is between 1 and this maximum sub-group size, and the size of an individual sub-group is invariant for the duration of a kernel's execution. Similarly to work-groups, the work-items within the same sub-group can be synchronized by calling the group_barrier function.

To maximize portability across devices, developers should not assume that work-items within a sub-group execute in any particular order, that work-groups are subdivided into sub-groups in a specific way, nor that two sub-groups within a work-group will make independent forward progress with respect to one another.

Variables with reduction semantics can be added to work-group data parallel kernels using the features described in Section 4.10.2.

3.8.5 Hierarchical data parallel kernels

[Note for this provisional version: Based on developer and implementation feedback, the hierarchical data parallel kernel feature described next is undergoing improvements to better align with the frameworks and patterns prevalent in modern programming. As this is a key part of the SYCL API and we expect to make changes to it, we temporarily recommend that new codes refrain from using this feature until the new API is finished in a near-future version of the SYCL specification, when full use of the updated feature will be recommended for use in new code. Existing codes using this feature will of course be supported by conformant implementations of this specification. — end note]

The SYCL compiler provides a way of specifying data parallel kernels that execute within work-groups via a different syntax which highlights the hierarchical nature of the parallelism. This mode is purely a compiler feature and does not change the execution model of the kernel. Instead of calling sycl::parallel_for the user calls sycl::parallel_for_work_group with a sycl::range value representing the number of work-groups to launch and optionally a second sycl::range representing the size of each work-group for performance tuning. All code within the parallel_for_work_group scope effectively executes once per work-group. Within the parallel_for_work_group scope, it is possible to call parallel_for_work_item which creates a new scope in which all work-items within the current work-group execute. This enables a programmer to write code that looks like there is an inner work-item loop inside an outer work-group loop, which closely matches the effect of the execution model. All variables declared inside the parallel_for_work_group scope are allocated in work-

group local memory, whereas all variables declared inside the parallel_for_work_item scope are declared in private memory. All parallel_for_work_item calls within a given parallel_for_work_group execution must have the same dimensions.

3.8.6 Kernels that are not launched over parallel instances

Simple kernels for which only a single instance of the kernel function will be executed are enqueued with the sycl ::single_task function. The kernel enqueued takes no "work-item id" parameter and will only execute once. The behavior is logically equivalent to executing a kernel on a single compute unit with a single work-group comprising only one work-item. Such kernels may be enqueued on multiple queues and devices and as a result may be executed in task-parallel fashion.

3.8.7 Pre-defined kernels

Some SYCL backends may expose pre-defined functionality to users as kernels. These kernels are not programmable, hence they are not bound by the SYCL C++ programming model restrictions, and how they are written is implementation-defined.

3.8.8 Synchronization

Synchronization of processing elements executing inside a device is handled by the SYCL device kernel following the SYCL kernel execution model. The synchronization of the different SYCL device kernels executing with the host memory is handled by the SYCL application via the SYCL runtime.

3.8.8.1 Synchronization in the SYCL application

Synchronization points between host and device(s) are exposed through the following operations:

• *Buffer destruction*: The destructors for sycl::buffer, sycl::unsampled_image and sycl::sampled_image objects wait for all submitted work on those objects to complete and to copy the data back to host memory before returning. These destructors only wait if the object was constructed with attached host memory and if data needs to be copied back to the host.

More complex forms of synchronization on buffer destruction can be specified by the user by constructing buffers with other kinds of references to memory, such as shared_ptr and unique_ptr.

- *Host Accessors*: The constructor for a host accessor waits for all kernels that modify the same buffer (or image) in any queues to complete and then copies data back to host memory before the constructor returns. Any command groups with requirements to the same memory object cannot execute until the host accessor is destroyed (see 3.5).
- Command group enqueue: The SYCL runtime internally ensures that any command groups added to queues have the correct event dependencies added to those queues to ensure correct operation. Adding command groups to queues never blocks. Instead any required synchronization is added to the queue and events of type sycl::event are returned by the queue's submit function that contain event information related to the specific command group.
- *Queue operations*: The user can manually use queue operations, such as **wait** to block execution of the calling thread until all the command groups submitted to the queue have finished execution. Note that this will also affect the dependencies of those command groups in other queues.

• *SYCL event objects*: SYCL provides sycl::event objects which can be used for synchronization. If synchronization is required across SYCL contexts from different SYCL backends, then the SYCL runtime ensures that extra host-based synchronization is added to enable the SYCL event objects to operate between contexts correctly.

Note that the destructors of other SYCL objects (sycl::queue, sycl::context,...) do not block. Only a sycl:: buffer, sycl::sampled_image or sycl::unsampled_image destructor might block. The rationale is that an object without any side effect on the host does not need to block on destruction as it would impact the performance. So it is up to the programmer to use a member function to wait for completion in some cases if this does not fit the goal. See Section 3.8.12 for more information on object life time.

3.8.8.2 Synchronization in SYCL kernels

In SYCL, synchronization can be either global or local within a group of work-items. Synchronization between work-items in a single group is achieved using a group barrier.

All the work-items of a group must execute the barrier before any are allowed to continue execution beyond the barrier. Note that the group barrier must be encountered by all work-items of a group executing the kernel or by none at all. In SYCL, work-group barrier and sub-group barrier functionality is exposed via the group_barrier function.

There is no mechanism for synchronization between work-items of different work-groups.

3.8.9 Error handling

In SYCL, there are two types of errors: synchronous errors that can be detected immediately when an API call is made, and asynchronous errors that can only be detected later after an API call has returned. Synchronous errors, such as failure to construct an object, are reported immediately by the runtime throwing an exception. Asynchronous errors, such as an error occurring during execution of a kernel on a device, are reported via an asynchronous error-handler mechanism.

Asynchronous errors are not reported immediately as they occur. The asynchronous error handler for a context or queue is called with a sycl::exception_list object, which contains a list of asynchronously-generated exception objects, on the conditions described by 4.15.1.1 and 4.15.1.2.

Asynchronous errors may be generated regardless of whether the user has specified any asynchronous error handler(s), as described in 4.15.1.2.

Some SYCL backends can report errors that are specific to the platform they are targeting, or that are more concrete than the errors provided by the SYCL API. Any error reported by a SYCL backend must derive from the base sycl::exception. When a user wishes to capture specifically an error thrown by a SYCL backend, she must include the SYCL backend-specific headers for said SYCL backend.

3.8.10 Fallback mechanism

A command group function object can be submitted either to a single queue to be executed on, or to a secondary queue. If a command group function object fails to be enqueued to the primary queue, then the system will attempt to enqueue it to the secondary queue, if given as a parameter to the submit function. If the command group function object fails to be queued to both of these queues, then a synchronous SYCL exception will be thrown.

It is possible that a command group may be successfully enqueued, but then asynchronously fail to run, for some

reason. In this case, it may be possible for the runtime system to execute the command group function object on the secondary queue, instead of the primary queue. The situations where a SYCL runtime may be able to achieve this asynchronous fall-back is implementation-defined.

3.8.11 Scheduling of kernels and data movement

A command group function object takes a reference to a command group handler as a parameter and anything within that scope is immediately executed and takes the handler object as a parameter. The intention is that a user will perform calls to SYCL functions, member functions, destructors and constructors inside that scope. These calls will be non-blocking on the host, but enqueue operations to the queue that the command group is submitted to. All user functions within the command group scope will be called on the host as the command group function object is executed, but any device kernels it invokes will be added to the SYCL queue. All kernels added to the queue will be executed out-of-order from each other, according to their data dependencies.

3.8.12 Managing object lifetimes

A SYCL application does not initialize any SYCL backend features until a sycl::context object is created. A user does not need to explicitly create a sycl::context object, but they do need to explicitly create a sycl::queue object, for which a sycl::context object will be implicitly created if not provided by the user.

All SYCL backend objects encapsulated in SYCL objects are reference-counted and will be destroyed once all references have been released. This means that a user needs only create a SYCL queue (which will automatically create an SYCL context) for the lifetime of their application to initialize and release any SYCL backend objects safely.

There is no global state specified to be required in SYCL implementations. This means, for example, that if the user creates two queues without explicitly constructing a common context, then a SYCL implementation does not have to create a shared context for the two queues. Implementations are free to share or cache state globally for performance, but it is not required.

Memory objects can be constructed with or without attached host memory. If no host memory is attached at the point of construction, then destruction of that memory object is non-blocking. The user may use C++ standard pointer classes for sharing the host data with the user application and for defining blocking, or non-blocking behavior of the buffers and images. If host memory is attached by using a raw pointer, then the default behavior is followed, which is that the destructor will block until any command groups operating on the memory object have completed, then, if the contents of the memory object is modified on a device those contents are copied back to host and only then does the destructor return. Instead of a raw pointer, a unique_ptr may be provided, which uses move semantics for initializing and using the associated host memory. In this case, the behavior of the buffer in relation to the user application and the SYCL runtime, then the reference counter of the shared_ptr determines whether the buffer needs to copy data back on destruction, and in that case the blocking or non-blocking behavior depends on the user application.

As said in Section 3.8.8, the only blocking operations in SYCL (apart from explicit wait operations) are:

- host accessor constructor, which waits for any kernels enqueued before its creation that write to the corresponding object to finish and be copied back to host memory before it starts processing. The host accessor does not necessarily copy back to the same host memory as initially given by the user;
- memory object destruction, in the case where copies back to host memory have to be done or when the host memory is used as a backing-store.

3.8.13 Device discovery and selection

A user specifies which queue to submit a command group function object and each queue is targeted to run on a specific device (and context). A user can specify the actual device on queue creation, or they can specify a device selector which causes the SYCL runtime to choose a device based on the user's provided preferences. Specifying a device selector causes the SYCL runtime to perform device discovery. No device discovery is performed until a SYCL device selector is passed to a queue constructor. Device topology may be cached by the SYCL runtime, but this is not required.

Device discovery will return all devices from all platforms exposed by all the supported SYCL backends, including the host backend.

3.8.14 Interfacing with SYCL backend API

There are two styles of developing a SYCL application: (1) Writing a pure SYCL generic application or (2) Writing a SYCL application that relies on some SYCL backend specific behavior.

When users follow (1), there is no assumption about what SYCL backend will be used during compilation or execution of the SYCL application. Therefore, the SYCL backend API is not assumed to be available to the developer. Only standard C++ types and interfaces are assumed to be available, as described in Section 3.8. Users only need to include the SYCL/sycl.hpp header to write a SYCL generic application.

On the other hand, when users follow (2), they must know what SYCL backend APIs they are using. In this case, any header required for the normal programmability of the SYCL backend API is assumed to be available to the user. In addition to the SYCL/sycl.hpp header, users must also include the SYCL backend-specific header as defined in Section 4.3. The SYCL backend-specific header provides the interoperability interface for the SYCL API to interact with native backend objects. Any type or header from the underlying SYCL backend API is included by the SYCL backend-specific header, under a namespace named after the backend name, e.g. namespace opencl would encapsulate the OpenCL headers. This avoids accidental pollution of user space with SYCL backend-specific types.

The interoperability API is defined in Section 4.5.2.

3.9 Memory objects

SYCL memory objects represent data that is handled by the SYCL runtime and can represent allocations in one or multiple devices at any time. Memory objects, both buffers and images, may have one or more underlying native backend objects to ensure that queues objects can use data in any device. A SYCL implementation may have multiple native backend objects for the same device. The SYCL runtime is responsible for ensuring the different copies are up-to-date whenever necessary, using whatever mechanism is available in the system to update the copies of the underlying native backend objects.

[Implementation note: A valid mechanism for this update is to transfer the data from one SYCL backend into the system memory using the SYCL backend-specific mechanism available, and then transfer it to a different device using the mechanism exposed by the new SYCL backend. — end note]

Memory objects in SYCL fall into one of two categories: buffer objects and image objects. A buffer object stores a one-, two- or three-dimensional collection of elements that are stored linearly directly back to back in the same way C or C++ stores arrays. An image object is used to store a one-, two- or three-dimensional texture, frame-buffer or image data that may be stored in an optimized and device-specific format in memory and must be accessed through specialized operations.

Elements of a buffer object can be a scalar data type (such as an int or float), vector data type, or a user-defined structure. In SYCL, a buffer object is a templated type (sycl::buffer), parameterized by the element type and number of dimensions. An image object is stored in one of a limited number of formats. The elements of an image object are selected from a list of predefined image formats which are provided by an underlying SYCL backend implementation. Images are encapsulated in the sycl::unsampled_image or sycl::sampled_image types, which are templated by the number of dimensions in the image. The minimum number of elements in a memory object is one.

The fundamental differences between a buffer and an image object are:

- Elements in a buffer are stored in an array of 1, 2 or 3 dimensions and can be accessed using an accessor by a kernel executing on a device. The accessors for kernels provide a member function to get C++ pointer types, or the sycl::global_ptr or sycl::constant_ptr classes.
- Elements of an image are stored in a format that is opaque to the user and cannot be directly accessed using a pointer. SYCL provides image accessors and samplers to allow a kernel to read from or write to an image.
- For a buffer object the data is accessed within a kernel in the same format as it is stored in memory, but in the case of an image object the data is not necessarily accessed within a kernel in the same format as it is stored in memory.
- Image elements are always a 4-component vector (each component can be a float or signed/unsigned integer) in a kernel. The SYCL accessor and sampler member functions to read from an image convert an image element from the format that it is stored into a 4-component vector.

Similarly, the SYCL accessor member functions provided to write to an image convert the image element from a 4-component vector to the appropriate image format specified such as four 8-bit elements, for example.

Users may want fine-grained control of the synchronization, memory management and storage semantics of SYCL image or buffer objects. For example, a user may wish to specify the host memory for a memory object to use, but may not want the memory object to block on destruction.

Depending on the control and the use cases of the SYCL applications, well established C++ classes and patterns can be used for reference counting and sharing data between user applications and the SYCL runtime. For control over memory allocation on the host and mapping between host and device memory, pre-defined or user-defined C++ std::allocator classes are used. For better control of synchronization between a SYCL and a non SYCL application that share data, std::shared_ptr and std::mutex classes are used.

3.10 SYCL device compiler

To enable SYCL to work on a variety of platforms, with different devices, operating systems, build systems and host compilers, SYCL provides a number of options to implementers to enable the compilation of SYCL kernels for devices, while still providing a unified programming model to the user.

3.10.1 Building a SYCL program

A SYCL program runs on a *host* and one or more SYCL devices. This requires a compilation model that enables compilation for a variety of targets. There is only ever one host for the SYCL program, so the compilation of the source code for the host must happen once and only once. Both kernel and non-kernel source code is compiled for the host.

The design of SYCL enables a single SYCL source file to be passed to multiple, different compilers, using the SMCP technique. This is an implementation option and is not required. What this option enables is for an implementer to provide a device compiler only and not have to provide a host compiler. A programmer who uses such an implementation will compile the same source file twice: once with the host compiler of their choice and once with a device compiler. This approach allows the advantages of having a single source file for both host code and kernels, while still allowing users an independent choice of host and SYCL device compilers.

Only the kernels are compiled for SYCL devices. Therefore, any compiler that compiles only for one or more devices must not compile non-kernel source code. Kernels are contained within C++ source code and may be dependent on lambda capture and template parameters, so compilation of the non-kernel code must determine lambda captures and template parameters, but not generate device code for non-kernel code.

Compilation of a SYCL program may follow either of the following options. The choice of option is made by the implementer:

- Separate compilation: One or more device compilers compile just the SYCL kernels for one or more devices. The device compilers all produce header files for interfacing between the host compiler and the SYCL runtime, which are integrated together with a tool that produces a single header file. The user compiles the source file with a normal C++ host compiler for their platform. The user must ensure that the host compiler is given the correct command-line arguments to ensure that the device compiler output header file is #included from inside the SYCL header files.
- 2. *Single-source compiler*: In this approach, a single compiler may compile an entire source file for both host and one or more devices. It is the responsibility of the single-source compiler to enable kernels to be compiled correctly for devices and enqueued from the host.

An implementer of SYCL may choose an implementation approach from the options above.

3.10.2 Naming of kernels

SYCL kernels are extracted from C++ source files and stored in an implementation-defined format. When the SYCL runtime needs to enqueue a SYCL kernel, it is necessary for the SYCL runtime to load the kernel and pass it to a SYCL backend API. This requires the kernel to have a name that is unique at enclosing namespace scope, to enable an association between the kernel invocation and the kernel itself. The association is achieved using a kernel name, which is a C++ type name.

For a named function object, the kernel name can be the same type as the function object itself, as long as the function object type is unique across the enclosing namespace scopes. For a lambda function the user may optionally provide a name for debugging or other reasons. In SYCL, this optional name is provided as a template parameter to the kernel invocation, e.g. parallel_for<class kernelName>, and this name may optionally be forward declared at namespace scope (but must always avoid conflict with another name at enclosing namespace scope).

A device compiler should detect the kernel invocations (e.g. parallel_for) in the source code and compile the enclosed kernels, storing them with their associated type name. For details please refer to 5.2.

3.11 Language restrictions in kernels

The SYCL kernels are executed on SYCL devices and all of the functions called from a SYCL kernel are going to be compiled for the device by a SYCL device compiler. Due to restrictions of the heterogeneous devices where the SYCL kernel will execute, there are certain restrictions on the base C++ language features that can be used

inside kernel code. For details on language restrictions please refer to 5.4.

SYCL kernels use parameters that are captured by value in the command group scope or are passed from the host to the device using the data management runtime classes of sycl::accessors. Sharing data structures between host and device code imposes certain restrictions, such as use of only user defined classes that are *C++ trivially copyable* classes for the data structures, and in general, no pointers initialized for the host can be used on the device. The only way of passing pointers to a kernel is through the sycl::accessor class, which supports the sycl::buffer, sycl::unsampled_image and sycl::sampled_image classes. No hierarchical structures of these classes are supported and any other data containers need to be converted to the SYCL data management classes using the SYCL interface. For more details on the rules for kernel parameter passing, please refer to 4.14.4.

3.11.1 SYCL linker

In SYCL only offline linking is supported for SYCL programs and libraries, however the mechanism is optional. In the case of linking C++ functions to a SYCL application, where the definitions are not available in the same translation unit of the compiler, then the macro SYCL_EXTERNAL has to be provided.

3.11.2 Functions and data types available in kernels

Inside kernels, the functions and data types available are restricted by the underlying capabilities of SYCL backend devices.

3.12 Endianness support

SYCL supports both big-endian and little-endian systems as long as it is supported by the used SYCL backends. However SYCL does not support mix-endian systems and does not support specifying the endianness of data within a SYCL kernel function.

Users must be aware of the endianness of the host and the SYCL backend devices they are targeting to ensure kernel arguments are processed correctly when applicable.

3.13 Example SYCL application

Below is a more complex example application, combining some of the features described above.

```
1
    #include <SYCL/sycl.hpp>
2
    #include <iostream>
3
4
   using namespace sycl;
5
   // Size of the matrices
6
7
    const size_t N = 2000;
8
    const size_t M = 3000;
9
10
   int main() {
     // Create a queue to work on
11
12
     queue myQueue;
13
14
      // Create some 2D buffers of float for our matrices
15
      buffer<float, 2> a { range<2>{N, M} };
```

```
16
      buffer<float, 2> b { range<2>{N, M} };
17
      buffer<float, 2> c { range<2>{N, M} };
18
19
      // Launch an asynchronous kernel to initialize a
20
      myQueue.submit([&](handler& cgh) {
21
        // The kernel writes a, so get a write accessor on it
22
        accessor A { a, cgh, write_only };
23
24
        // Enqueue a parallel kernel iterating on a N*M 2D iteration space
25
        cgh.parallel_for(range<2> {N, M}, [=](id<2> index) {
26
          A[index] = index[0] * 2 + index[1];
27
        });
28
      });
29
30
      // Launch an asynchronous kernel to initialize b
31
      myQueue.submit([&](handler& cgh) {
32
        // The kernel writes b, so get a write accessor on it
33
        accessor B { b, cgh, write_only };
34
35
        // From the access pattern above, the SYCL runtime detects that this
        // command_group is independent from the first one and can be
36
37
        // scheduled independently
38
39
        // Enqueue a parallel kernel iterating on a N*M 2D iteration space
40
        cgh.parallel_for(range<2> {N, M}, [=](id<2> index) {
41
          B[index] = index[0] * 2014 + index[1] * 42;
42
        });
43
      });
44
45
      // Launch an asynchronous kernel to compute matrix addition c = a + b
46
      myQueue.submit([&](handler& cgh) {
47
        // In the kernel a and b are read, but c is written
48
        accessor A { a, cgh, read_only };
49
        accessor B { b, cgh, read_only };
50
        accessor C { c, cgh, write_only };
51
52
        // From these accessors, the SYCL runtime will ensure that when
53
        /\!/ this kernel is run, the kernels computing a and b have completed
54
55
        // Enqueue a parallel kernel iterating on a N*M 2D iteration space
56
        cgh.parallel_for(range<2> {N, M}, [=](id<2> index) {
57
            C[index] = A[index] + B[index];
58
         });
59
        });
60
61
      // Ask for an accessor to read c from application scope. The SYCL runtime
      // waits for c to be ready before returning from the constructor
62
63
      host_accessor C { c, read_only };
      std::cout << std::endl << "Result:" << std::endl;</pre>
64
65
      for (size_t i = 0; i < N; i++) {</pre>
        for (size_t j = 0; j < M; j++) {
66
67
          // Compare the result to the analytic value
68
          if (C[i][j] != i * (2 + 2014) + j * (1 + 42)) {
69
            std::cout << "Wrong value " << C[i][j] << " on element " << i << " "</pre>
70
                      << j << std::endl;
```

71 exit(-1);
72 }
73 }
74 }
75
76 std::cout << "Good computation!" << std::endl;
77 return 0;
78 }</pre>

4. SYCL programming interface

The SYCL programming interface provides a common abstracted feature set to one or more SYCL backend APIs. This section describes the C++ library interface to the SYCL runtime which executes across those SYCL backends.

The entirety of the SYCL interface defined in this section is required to be available for any SYCL backends, with the exception of the interoperability interface, which is described in general terms in this document, not pertaining to any particular SYCL backend.

SYCL guarantees that all the member functions and special member functions of the SYCL classes described are thread safe.

4.1 Backends

The SYCL backends that are available to a SYCL implementation can be identified using the enum class backend

```
1 namespace sycl {
2 enum class backend {
3 <see-below>
4 };
5 } // namespace sycl
```

The enum class backend is implementation defined and must be populated with a unique identifier for each SYCL backend that the SYCL implementation supports, containing at least one SYCL backend that is a SYCL host backend.

Each named SYCL backend enumerated in the enum class backend must be associated with a SYCL backend specification. Many sections of this specification will refer to the associated SYCL backend specification.

4.1.1 Backend macros

As the identifiers defined in enum class backend are implementation defined a SYCL implementation must also define a pre-processor macro for each of these identifiers. If the SYCL backend is defined by the Khronos SYCL group, the name of the macro has the form SYCL_BACKEND_<backend_name>, where *backend_name* is the associated identifier from backend in all upper-case. See Chapter 6 for the name of the macro if the vendor defines the SYCL backend outside of the Khronos SYCL group.

4.2 Generic vs non-generic SYCL

The SYCL programming API is split into two categories; generic SYCL and non-generic SYCL. Almost everything in the SYCL programming API is considered generic SYCL. However any usage of the enum class

backend is considered non-generic SYCL and should only be used for SYCL backend specialized code paths, as the identifiers defined in backend are implementation defined.

In any non-generic SYCL application code where the backend enum class is used, the expression must be guarded with a pre-process **#ifdef** guard using the associated pre-process macro to ensure that the SYCL application will compile even if the SYCL implementation does not support that SYCL backend being specialized for.

4.3 Header files and namespaces

SYCL provides one standard header file: "SYCL/sycl.hpp", which needs to be included in every translation unit which uses the SYCL programming API.

All SYCL classes, constants, types and functions defined by this specification should exist within the ::sycl namespace.

For compatibility with SYCL 1.2.1, SYCL provides another standard header file: "CL/sycl.hpp", which can be included in place of "SYCL/sycl.hpp".

In that case, all SYCL classes, constants, types and functions defined by this specification should exist within the ::cl::sycl C++ namespace.

For consistency, the programming API will only refer to the "sycl.hpp" header and the ::sycl namespace, but this should be considered synonymous with the SYCL 1.2.1 header and namespace.

The sycl::detail namespace is reserved for implementation details.

When a SYCL backend is defined by the Khronos SYCL group, functionality for that SYCL backend is available via the header "SYCL/backend/<backend_name>.hpp", and all SYCL backend-specific functionality is made available in the namespace sycl::<backend_name> where *backend_name* is the name of the SYCL backend as defined in the SYCL backend specification.

Chapter 6 defines the allowable header files and namespaces for any extensions that a vendor may provide, including any SYCL backend that the vendor may define outside of the Khronos SYCL group.

4.4 Class availability

In SYCL some SYCL runtime classes are available to the SYCL application, some are available within a SYCL kernel function and some are available on both and can be passed as parameters to a SYCL kernel function.

Each of the following SYCL runtime classes: buffer, buffer_allocator, context, device, event, exception , handler, id, image_allocator, kernel, marray, module, nd_range, platform, queue, range, sampled_image, sampler, stream, unsampled_image and vec must be available to the host application.

Each of the following SYCL runtime classes: accessor, atomic_ref, device_event, group, h_item, id, item, marray, multi_ptr, nd_item, range, reducer, sampler, stream, sub_group and vec must be available within a SYCL kernel function.

Each of the following SYCL runtime classes: accessor, id, marray, range, reducer, sampler, stream and vec are permitted as parameters to a SYCL kernel function.

4.5 Common interface

When a dimension template parameter is used in SYCL classes, it is defaulted as 1 in most cases.

4.5.1 Param traits class

The class param_traits is a C++ type trait for providing an alias to the return type associated with each info parameter. An implementation must provide a specialization of the param_traits class for every info parameter with the associated return type as defined in the info parameter tables.

```
1
   namespace sycl {
2
    namespace info {
3
    template <typename T, T param>
4
    class param_traits {
5
    public:
6
7
       using return_type = __return_type__<T, param>;
8
9
   };
10 } // namespace info
11 } // namespace sycl
```

4.5.2 Backend interoperability

Many of the SYCL runtime classes may be implemented such that they encapsulate an object unique to the SYCL backend that underpins the functionality of that class. Where appropriate, these classes may provide an interface for interoperating between the SYCL runtime object and the native backend object in order to support interoperability within an applications between SYCL and the associated SYCL backend API.

There are two forms of interoperability with SYCL runtime classes; interoperability on the SYCL application with the SYCL backend API and interoperability within a SYCL kernel function with the equivalent kernel language types of the SYCL backend. SYCL application interoperability and SYCL kernel function interoperability are provided via different interfaces and may have different native backend object types.

SYCL application interoperability may be provided for buffer, context, device, event, kernel, module, platform, queue, sampled_image, image_sampler, stream and unsampled_image.

SYCL kernel function interoperability may be provided for accessor, stream and device_event inside the SYCL kernel function scope only. SYCL kernel function interoperability is not available inside command group scope.

Support for SYCL backend interoperability is optional and therefore not required to be provided by a SYCL implementation. A SYCL application using SYCL backend interoperability is considered to be non-generic SYCL.

Details on the interoperability for a given SYCL backend are available on the SYCL backend specification document for that SYCL backend.

4.5.2.1 Type traits backend_traits

```
1 namespace sycl {
2
```

```
3 template <backend Backend>
    class backend traits {
 4
 5
     public:
 6
 7
      template <class T>
 8
      using native_type = see-below;
 9
10
      using errc = see-below;
11
12 };
13
14 } // namespace sycl
```

A series of type traits are provided for SYCL backend interoperability, defined in the backend_traits class.

A specialization of **backend_traits** must be provided for each named **SYCL** backend enumerated in the enum class backend.

- For each SYCL runtime class T which supports SYCL application interoperability with the SYCL backend, a specialisation of native_type must be defined as the type of SYCL application interoperability native backend object associated with T for the SYCL backend, specified in the SYCL backend specification.
- For each SYCL runtime class T which supports kernel function interoperability with the SYCL backend, a specialisation of native_type within backend_traits must be defined as the type of the kernel function interoperability native backend object associated with T for the SYCL backend, specified in the backend specification.
- A specialization of errc must be defined as the SYCL backend error code type.

4.5.2.2 Template function get_native

```
1 namespace sycl {
2
3 template<backend Backend, class T>
4 backend_traits<Backend>::native_type<T> get_native(const T &syclObject);
5
6 } // namespace sycl
```

For each SYCL runtime class T which supports SYCL application interoperability, a specialisation of get_native must be defined, which takes an instance of T and returns a SYCL application interoperability native backend object associated with syclObject which can be used for SYCL application interoperability. The lifetime of the object returned are backend-defined and specified in the backend specification.

For each SYCL runtime class T which supports kernel function interoperability, a specialisation of get_native must be defined, which takes an instance of T and returns the kernel function interoperability native backend object associated with syclObject which can be used for kernel function interoperability. The lifetime of the object returned are backend-defined and specified in the backend specification.

4.5.2.3 Template functions make_*

```
1 namespace sycl {
2
```

```
3 template<backend Backend>
 4
    platform make_platform(const backend_traits<Backend>::native_type<platform> &backendObject);
 5
 6 template<backend Backend>
 7
   device make_device(const backend_traits<Backend>::native_type<device> &backendObject);
 8
 9 template<backend Backend>
10 context make_context(const backend_traits<Backend>::native_type<context> &backendObject, const
        async_handler asyncHandler = {});
11
12 template<backend Backend>
    queue make_queue(const backend_traits<Backend>::native_type<queue> &backendObject,
13
14
                     const context &targetContext, const async_handler asyncHandler = {});
15
16 template<backend Backend>
    event make_event(const backend_traits<Backend>::native_type<event> &backendObject,
17
18
                     const context &targetContext);
19
20 template<backend Backend>
21
    buffer make_buffer(const backend_traits<Backend>::native_type<buffer> &backendObject,
22
                       const context &targetContext, event availableEvent = {});
23
24 template<backend Backend>
25
    sampled_image make_sampled_image(
26
      const backend_traits<Backend>::native_type<sampled_image> &backendObject,
27
      const context &targetContext, image_sampler imageSampler, event availableEvent = {});
28
29 template<backend Backend>
30
   unsampled_image make_unsampled_image(
31
      const backend_traits<Backend>::native_type<unsampled_image> &backendObject,
32
      const context &targetContext, event availableEvent = {});
33
34 template<backend Backend>
35 image_sampler make_image_sampler(
36
      const backend_traits<Backend>::native_type<image_sampler> &backendObject,
37
     const context &targetContext);
38
39 template<backend Backend>
40
    stream make_stream(const backend_traits<Backend>::native_type<stream> &backendObject,
41
                       const context &targetContext, event availableEvent = {});
42
43 template<backend Backend>
44 kernel make_kernel(const backend_traits<Backend>::native_type<kernel> &backendObject,
45
                       const context &targetContext);
46
47 template<backend Backend>
48 kernel make_module(const backend_traits<Backend>::native_type<event> &backendObject,
49
                       const context &targetContext);
50
51 } // namespace sycl
```

For each SYCL runtime class T which supports SYCL application interoperability, a specialisation of the appropriate template function make_{sycl_class} where {sycl_class} is the class name of T, must be defined, which takes a SYCL application interoperability native backend object and constructs and returns an instance of T. The lifetime of the object returned is backend-defined and specified in the backend specification.

4.5.3 Common reference semantics

Each of the following SYCL runtime classes: accessor, buffer, context, device, event, kernel, module, platform, queue, sampled_image, sampler and unsampled_image, must obey the following statements, where T is the runtime class type:

- T must be copy constructible and copy assignable on the host application and within SYCL kernel functions in the case that T is a valid kernel argument. Any instance of T that is constructed as a copy of another instance, via either the copy constructor or copy assignment operator, must behave as-if it were the original instance and as-if any action performed on it were also performed on the original instance and if said instance is not a host object must represent and continue to represent the same underlying native backend object as the original instance where applicable.
- T must be destructible on the host application and within SYCL kernel functions in the case that T is a valid kernel argument. When any instance of T is destroyed, including as a result of the copy assignment operator, any behavior specific to T that is specified as performed on destruction is only performed if this instance is the last remaining host copy, in accordance with the above definition of a copy and the destructor requirements described in 4.5.2 where applicable.
- T must be move constructible and move assignable on the host application and within SYCL kernel functions in the case that T is a valid kernel argument. Any instance of T that is constructed as a move of another instance, via either the move constructor or move assignment operator, must replace the original instance rendering said instance invalid and if said instance is not a host object must represent and continue to represent the same underlying native backend object as the original instance where applicable.
- T must be equality comparable on the host application. Equality between two instances of T (i.e. a == b) must be true if one instance is a copy of the other and non-equality between two instances of T (i.e. a != b) must be true if neither instance is a copy of the other, in accordance with the above definition of a copy, unless either instance has become invalidated by a move operation. By extension of the requirements above, equality on T must guarantee to be reflexive (i.e. a == a), symmetric (i.e. a == b implies b == a and a != b implies b != a) and transitive (i.e. a == b && b == c implies c == a).
- A specialization of std::hash for T must exist on the host application that returns a unique value such that if two instances of T are equal, in accordance with the above definition, then their resulting hash values are also equal and subsequently if two hash values are not equal, then their corresponding instances are also not equal, in accordance with the above definition.

Some SYCL runtime classes will have additional behavior associated with copy, movement, assignment or destruction semantics. If these are specified they are in addition to those specified above unless stated otherwise.

Each of the runtime classes mentioned above must provide a common interface of special member functions in order to fulfill the copy, move, destruction requirements and hidden friend functions in order to fulfill the equality requirements.

A hidden friend function is a function first declared via a **friend** declaration with no additional out of class or namespace scope declarations. Hidden friend functions are only visible to ADL (Argument Dependent Lookup) and are hidden from qualified and unqualified lookup. Hidden friend functions have the benefits of avoiding accidental implicit conversions and faster compilation.

These common special member functions and hidden friend functions are described in Tables 4.1 and 4.2 respectively.

```
1
   namespace sycl {
 2
 3
   class T {
 4
      ...
 5
 6
     public:
 7
     T(const T &rhs);
 8
 9
     T(T &&rhs);
10
      T &operator=(const T &rhs);
11
12
      T &operator=(T &&rhs);
13
14
15
     ~T();
16
17
      . . .
18
      friend bool operator==(const T &lhs, const T &rhs) { /* ... */ }
19
20
      friend bool operator!=(const T &lhs, const T &rhs) { /* ... */ }
21
22
23
      . . .
24 };
25 } // namespace sycl
```

| Special member function | Description |
|----------------------------|---|
| T(const T &rhs) | Constructs a T instance as a copy of the |
| | RHS SYCL T in accordance with the re- |
| | quirements set out above. |
| T(T &&rhs) | Constructs a SYCL T instance as a move |
| | of the RHS SYCL T in accordance with the |
| | requirements set out above. |
| T &operator=(const T &rhs) | Assigns this SYCL T instance with a copy |
| | of the RHS SYCL T in accordance with the |
| | requirements set out above. |
| T &operator=(T &&rhs) | Assigns this SYCL T instance with a move |
| | of the RHS SYCL T in accordance with the |
| | requirements set out above. |
| ~T() | Destroys this SYCL T instance in accor- |
| | dance with the requirements set out in 4.5.3. |
| | On destruction of the last copy, may per- |
| | form additional lifetime related operations |
| | required for the underlying native backend |
| | object specified in the SYCL backend spec- |
| | ification, if this SYCL T instance was origi- |
| | nally constructed using one of the backend |
| | interoperability make_* functions specified |
| | in 4.5.2.3. |
| | End of table |

Table 4.1: Common special member functions for reference semantics.

| Hidden friend function | Description |
|--|---|
| <pre>bool operator==(const T &lhs, const T &rhs)</pre> | Returns true if this LHS SYCL T is equal to the RHS SYCL T in accordance with the re- |
| | false. |
| <pre>bool operator!=(const T &lhs, const T &rhs)</pre> | Returns true if this LHS SYCL T is not equal to the RHS SYCL T in accordance with the requirements set out above, other- wise returns false. |
| | End of table |

Table 4.2: Common hidden friend functions for reference semantics.

4.5.4 Common by-value semantics

Each of the following SYCL runtime classes: id, range, item, nd_item, h_item, group, sub_group and nd_range must follow the following statements, where T is the runtime class type:

- T must be default copy constructible and copy assignable on the host application (in the case where T is available on the host) and within SYCL kernel functions.
- T must be default destructible on the host application (in the case where T is available on the host) and within SYCL kernel functions.
- T must be default move constructible and default move assignable on the host application (in the case where T is available on the host) and within SYCL kernel functions.
- T must be equality comparable on the host application (in the case where T is available on the host) and within SYCL kernel functions. Equality between two instances of T (i.e. a == b) must be true if the value of all members are equal and non-equality between two instances of T (i.e. a != b) must be true if the value of any members are not equal, unless either instance has become invalidated by a move operation. By extension of the requirements above, equality on T must guarantee to be reflexive (i.e. a == a), symmetric (i.e. a == b implies b == a and a != b implies b != a) and transitive (i.e. a == b && b == c implies c == a).

Some SYCL runtime classes will have additional behavior associated with copy, movement, assignment or destruction semantics. If these are specified they are in addition to those specified above unless stated otherwise.

Each of the runtime classes mentioned above must provide a common interface of special member functions and member functions in order to fulfill the copy, move, destruction and equality requirements, following the rule of five and the rule of zero.

These common special member functions and hidden friend functions are described in Tables 4.3 and 4.4 respectively.

```
1 namespace sycl {
2
3 class T {
4 ...
5
6 public:
```

```
7
      // If any of the following five special member functions are not
 8
      // public, inline or defaulted, then all five of them should be
 9
      // explicitly declared (see rule of five).
10
      // Otherwise, none of them should be explicitly declared
11
      // (see rule of zero).
12
      // T(const T &rhs);
13
14
15
      // T(T &&rhs);
16
17
      // T &operator=(const T &rhs);
18
19
      // T &operator=(T &&rhs);
20
21
      // ~T();
22
23
      . . .
24
25
      friend bool operator==(const T &lhs, const T &rhs) { /* ... */ }
26
27
      friend bool operator!=(const T &lhs, const T &rhs) { /* ... */ }
28
29
      . . .
30 };
31 } // namespace sycl
```

| Special member function (see rule of five & rule of zero) | Description |
|---|---------------------------|
| T(const T &rhs); | Copy constructor. |
| T(T &&rhs); | Move constructor. |
| T &operator=(const T &rhs); | Copy assignment operator. |
| T &operator=(T &&rhs); | Move assignment operator. |
| ~T(); | Destructor. |
| | End of table |

Table 4.3: Common special member functions for by-value semantics.

| Hidden friend function | Description |
|--|---|
| <pre>bool operator==(const T &lhs, const T &rhs)</pre> | Returns true if this LHS SYCL T is equal to |
| | the RHS SYCL 1 in accordance with the re- |
| | quirements set out above, otherwise returns |
| | false. |
| <pre>bool operator!=(const T &lhs, const T &rhs)</pre> | Returns true if this LHS SYCL T is not |
| | equal to the RHS SYCL T in accordance |
| | with the requirements set out above, other- |
| | wise returns false. |
| | End of table |

Table 4.4: Common hidden friend functions for by-value semantics.

4.5.5 Properties

Each of the following SYCL runtime classes: accessor, context, queue, buffer, unsampled_image and sampled_image provide an optional parameter in each of their constructors to provide a property_list which contains zero or more properties. Each of those properties augments the semantics of the class with a particular feature. Each of those classes must also provide has_property and get_property member functions for querying for a particular property.

The listing below illustrates the usage of various buffer properties, described in 4.7.2.2.

The example illustrates how using properties does not affect the type of the object, thus, does not prevent the usage of SYCL objects in containers.

```
1
   {
 2
      context myContext;
 3
 4
      std::vector<buffer<int, 1>> bufferList {
 5
        buffer<int, 1>{ptr, rng},
 6
        buffer<int, 1>{ptr, rng, property::use_host_ptr{}},
 7
        buffer<int, 1>{ptr, rng, property::context_bound{myContext}}
 8
      };
 9
10
      for(auto& buf : bufferList) {
        if (buf.has_propertyproperty::context_bound>()) {
11
12
          auto prop = buf.get_propertyproperty::context_bound>();
13
          assert(myContext == prop.get_context());
14
        }
15
      }
16 }
```

Each property is represented by a unique class and an instance of a property is an instance of that type. Some properties can be default constructed while others will require an argument on construction. A property may be applicable to more than one class, however some properties may not be compatible with each other. See the requirements for the properties of the SYCL buffer class, SYCL unsampled_image class and SYCL sampled_image class in Table 4.33 and Table 4.40 respectively.

Any property that is provided to a SYCL runtime class via an instance of the SYCL property_list class must become encapsulated by that class and therefore shared between copies of that class. As a result properties must inherit the copy and move semantics of that class as described in 4.5.3.

A SYCL implementation or a SYCL backend may provide additional properties other than those defined here, provided they are defined in accordance with the requirements described in 4.3.

4.5.5.1 Properties interface

Each of the runtime classes mentioned above must provide a common interface of member functions in order to fulfill the property interface requirements.

A synopsis of the common properties interface, the SYCL property_list class and the SYCL property classes is provided below. The member functions of the common properties interface are listed in Table 4.6. The constructors of the SYCL property_list class are listed in Table 4.7.

```
1 namespace sycl {
2
3 template <typename propertyT>
4
   struct is_property;
5
6
   template <typename propertyT, typename syclObjectT>
7
   struct is_property_of;
8
9
   class T {
10
     . . .
11
12
     template <typename propertyT>
13
     bool has_property() const;
14
15
     template <typename propertyT>
16
     propertyT get_property() const;
17
18
     ...
19 };
20
21 class property_list {
22
    public:
23
      template <typename... propertyTN>
24
      property_list(propertyTN... props);
25 };
26 } // namespace sycl
```

| Traits | Description |
|--|---|
| <pre>template <typename propertyt=""></typename></pre> | An explicit specialization of is_property |
| <pre>struct is_property</pre> | that inherits from <pre>std::true_type</pre> must |
| | be provided for each property, where |
| | propertyT is the class defining the prop- |
| | erty. This includes both standard properties |
| | described in this specification and any ad- |
| | ditional non-standard properties defined by |
| | an implementation. All other specializa- |
| | tions of is_property must inherit from std |
| | ::false_type. |
| <pre>template <typename propertyt,="" syclobjectt=""></typename></pre> | An explicit specialization of |
| <pre>struct is_property_of</pre> | is_property_of that inherits from std |
| | ::true_type must be provided for each |
| | property that can be used in constructing a |
| | given SYCL class, where propertyT is the |
| | class defining the property and sycl0bjectT |
| | is the SYCL class. This includes both |
| | standard properties described in this spec- |
| | ification and any additional non-standard |
| | properties defined by an implementation. |
| | All other specializations of is_property_of |
| | must inherit from <pre>std::false_type.</pre> |
| | E 1 6 11 |

End of table

Table 4.5: Traits for properties.

| Member function | Description |
|--|--|
| <pre>template <typename propertyt=""></typename></pre> | Returns true if T was constructed with the |
| <pre>bool has_property()const</pre> | property specified by propertyT. Returns |
| | false if it was not. |
| <pre>template <typename propertyt=""></typename></pre> | Returns a copy of the property of type |
| <pre>propertyT get_property()const</pre> | propertyT that T was constructed with. |
| | Must throw an exception with the errc:: |
| | invalid_object_error error code if T was |
| | not constructed with the propertyT prop- |
| | erty. |
| | End of table |

Table 4.6: Common member functions of the SYCL property interface.

| Constructor | Description |
|---|--|
| <pre>template <typename propertytn=""></typename></pre> | Available only when: is_property< |
| <pre>property_list(propertyTN props)</pre> | <pre>property>::value evaluates to true where</pre> |
| | property is each property in propertyTN. |
| | Construct a SYCL property_list with zero |
| | or more properties. |
| | End of table |

Table 4.7: Constructors of the SYCL property_list class.

4.6 SYCL runtime classes

4.6.1 Device selection

Since a system can have several SYCL-compatible devices attached, it is useful to have a way to select a specific device or a set of devices to construct a specific object such as a device (see Section 4.6.4) or a queue (see Section 4.6.5), or perform some operations on a device subset.

Device selection is done either by having already a specific instance of a device (see Section 4.6.4) or by providing a device selector which is a ranking function that will give an integer ranking value to all the devices on the system.

4.6.1.1 Device selector

The actual interface for a device selector is a callable taking a const device reference and returning a value implicitly convertible to a int.

At any point where the SYCL runtime needs to select a SYCL device using a device selector, the system will query all available SYCL devices from all SYCL backends in the system including the SYCL host backend, will call the device selector on each device and select the one which returns the highest score. If the highest value is negative no device is selected.

In places where only one device has to be picked and the high score is obtained by more than one device, then one of the tied devices will be returned, but which one is not defined and may depend on enumeration order, for example, outside the control of the SYCL runtime.

| Some predefined device selectors are provided by the system | as described on | Table 4.8 in a | header file | with some |
|---|-----------------|------------------|-------------|-----------|
| definition similar to the following: | | | | |
| | | | | |

| SYCL device selectors | Description |
|------------------------|---|
| default_selector_v | Select a SYCL device from any sup- |
| | ported SYCL backend based on an |
| | implementation-defined heuristic. Must |
| | select the host device if no other suitable |
| | device can be found. |
| gpu_selector_v | Select a SYCL device from any sup- |
| | ported SYCL backend for which the de- |
| | <pre>vice type is info::device::device_type::</pre> |
| | gpu. The SYCL class constructor using |
| | it must throw an exception with the errc |
| | ::runtime_error error code if no device |
| | matching this requirement can be found. |
| accelerator_selector_v | Select a SYCL device from any sup- |
| | ported SYCL backend for which the de- |
| | <pre>vice type is info::device::device_type::</pre> |
| | accelerator. The SYCL class construc- |
| | tor using it must throw an exception with |
| | the errc::runtime_error error code if no |
| | device matching this requirement can be |
| | found. |
| cpu_selector_v | Select a SYCL device from any sup- |
| | ported SYCL backend for which the de- |
| | <pre>vice type is info::device::device_type::</pre> |
| | cpu. The SYCL class constructor using |
| | it must throw an exception with the errc |
| | ::runtime_error error code if no device |
| | matching this requirement can be found. |
| host_selector_v | Select the SYCL host device from the |
| | SYCL host backend. This must always re- |
| | turn a valid SYCL device. |
| | End of table |

Table 4.8: Standard device selectors included with all SYCL implementations.

1 namespace sycl { 2 3 // Predefined device selectors 4 __unspecified__ default_selector_v; __unspecified__ host_selector_v; 5 __unspecified__ cpu_selector_v; 6 7 __unspecified__ gpu_selector_v; 8 __unspecified__ accelerator_selector_v; 9 10 // Predefined types for compatibility with old SYCL 1.2.1 device selectors 11 using default_selector = __unspecified__; 12 using host_selector = __unspecified__;

```
13 using cpu_selector = __unspecified__;
14 using gpu_selector = __unspecified__;
15 using accelerator_selector = __unspecified__;
16
17 } // namespace sycl
```

Typical examples of default and user-provided device selectors could be:

```
sycl::device my_gpu { sycl::gpu_selector_v };
 1
2
3
   sycl::queue my_accelerator { sycl::accelerator_selector_v };
4
5 int prefer_my_vendor(const sycl::device & d) {
     // Return 1 if the vendor name is "MyVendor" or 0 else.
6
7
      // O does not prevent another device to be picked as a second choice
8
     return d.get_info<info::device::vendor>() == "MyVendor";
   }
9
10
11
   // Get the preferred device or another one if not available
   sycl::device preferred_device { prefer_my_vendor };
12
13
14 // This throws if there is no such device in the system
15 sycl::queue half_precision_controller {
     // Can use a lambda as a device ranking function.
16
17
     // Returns a negative number to fail in the case there is no such device
18
     [] (auto &d) { return d.has(aspect::fp16) ? 1 : -1; }
19 };
20
21
   // To ease porting SYCL 1.2.1 code, there are types whose
22
   // construction leads to the equivalent predefined device selector
23 sycl::queue my_old_style_gpu { sycl::gpu_selector {} };
```

[Note: in SYCL 1.2.1 the predefined device selectors were actually types that had to be instantiated to be used. Now they are just instances. To simplify porting code using the old type instantiations, a backward-compatible API is still provided, such as $sycl::default_selector$. The new predefined device selectors have their new names appended with _v to avoid conflicts, thus following the naming style used by traits in the C++ standard library. There is no requirement for the implementation to have for example $sycl::gpu_selector_v$ being an instance of $sycl::gpu_selector.$ — end note]

Implementation note: the SYCL API might rely on SFINAE or C++20 concepts to resolve some ambiguity in constructors with default parameters.

4.6.2 Platform class

The SYCL platform class encapsulates a single SYCL platform on which SYCL kernel functions may be executed. A SYCL platform must be associated with a single SYCL backend and may encapsulate a native backend object.

A SYCL platform is also associated with one or more SYCL devices associated with the same SYCL backend.

All member functions of the platform class are synchronous and errors are handled by throwing synchronous SYCL exceptions.

The default constructor of the SYCL platform class will construct a platform associated with the SYCL host backend.

The explicit constructor of the SYCL platform class which takes a device selector will construct a platform that is associated with the SYCL backend that is associated with the device that the device selector would construct, according to Sections 4.6.1.1 and 4.6.4.

The SYCL platform class provides the common reference semantics (see Section 4.5.3).

4.6.2.1 Platform interface

A synopsis of the SYCL platform class is provided below. The constructors, member functions and static member functions of the SYCL platform class are listed in Tables 4.9, 4.10 and 4.11 respectively. The additional common special member functions and common member functions are listed in 4.5.3 in Tables 4.1 and 4.2 respectively.

```
namespace sycl {
 1
    class platform {
2
3
    public:
 4
     platform();
 5
     template <typename DeviceSelector>
6
7
      explicit platform(const DeviceSelector &deviceSelector);
8
      /* -- common interface members -- */
9
10
11
     backend get_backend() const;
12
13
      std::vector<device> get_devices(
14
        info::device_type = info::device_type::all) const;
15
16
      template <info::platform param>
17
      typename info::param_traits<info::platform, param>::return_type get_info() const;
18
19
      template <typename BackendEnum, BackendEnum param>
20
      typename info::param_traits<BackendEnum, param>::return_type
21
      get_backend_info() const;
22
23
     bool has(aspect asp) const;
24
25
     bool has_extension(const std::string &extension) const; // Deprecated
26
27
     bool is_host() const;
28
29
     static std::vector<platform> get_platforms();
30 };
31 } // namespace sycl
```

| Constructor | Description |
|-------------|---|
| platform() | Constructs a SYCL platform instance as a host platform. |
| | Continued on next page |

Table 4.9: Constructors of the SYCL platform class.

| Constructor | Description |
|---|---|
| <pre>template <typename deviceselector=""></typename></pre> | Constructs a SYCL platform instance using the device se- |
| <pre>explicit platform(const DeviceSelector</pre> | lector parameter. One of the SYCL devices that is associ- |
| &) | ated with the constructed SYCL platform instance must be |
| | the SYCL device that is produced from the provided device |
| | ranking function. |
| | End of table |

| Member function | Description |
|--|--|
| <pre>backend get_backend()const</pre> | Returns a backend identifying the SYCL |
| | backend associated with this platform. |
| <pre>template <info::platform param=""></info::platform></pre> | Queries this SYCL platform for infor- |
| <pre>typename info::param_traits<info::platform,< pre=""></info::platform,<></pre> | mation requested by the template param- |
| param>::return_type | eter param. Specializations of info:: |
| <pre>get_info()const</pre> | param_traits must be defined in accor- |
| | dance with the info parameters in Table 4.19 |
| | to facilitate returning the type associated |
| | with the param parameter. |
| template <typename backendenum="" backendenum,="" param=""></typename> | Queries this SYCL platform for SYCL |
| <pre>typename info::param_traits<backendenum, param="">::</backendenum,></pre> | backend-specific information requested by |
| return_type | the template parameter param. BackendEnum |
| <pre>get_backend_info()const</pre> | can be any enum class type specified |
| | by the SYCL backend specification of a |
| | supported SYCL backend named accord- |
| | ing to the convention info:: <backend_name< td=""></backend_name<> |
| | >::platform and param must be a valid |
| | enumeration of that enum class. Spe- |
| | cializations of info::param_traits must |
| | be defined for BackendEnum in accor- |
| | dance with the SYCL backend specification. |
| | Must throw an exception with the errc:: |
| | invalid_object_error if the SYCL back- |
| | end that corresponds with BackendEnum is |
| | different from the SYCL backend that is as- |
| | sociated with this platform. |
| <pre>bool has(aspect asp)const</pre> | Returns true if all of the SYCL devices as- |
| | sociated with this SYCL platform have the |
| | given aspect. |
| <pre>bool has_extension(const std::string & extension)</pre> | Deprecated, use has () instead. |
| const | Returns true if this SYCL platform supports |
| | the extension queried by the extension pa- |
| | rameter. A SYCL platform can only sup- |
| | port an extension if all associated SYCL |
| | devices support that extension. |

Table 4.9: Constructors of the SYCL platform class.

Table 4.10: Member functions of the SYCL platform class.

Continued on next page

| Member function | Description |
|---|--|
| <pre>bool is_host()const</pre> | Returns true if the backend associated with |
| | this SYCL platform is a SYCL host back- |
| | end. |
| <pre>std::vector<device> get_devices(</device></pre> | Returns a std::vector containing all |
| <pre>info::device_type = info::device_type::all)const</pre> | SYCL devices associated with this SYCL |
| | <pre>platform. The returned std::vector must</pre> |
| | contain only a single SYCL device that |
| | is a host device if this SYCL platform |
| | is a host platform. Must return an empty |
| | std::vector instance if there are no devices |
| | that match the given info::device_type. |
| | End of table |

Table 4.10: Member functions of the SYCL platform class.

| Static member function | Description |
|--|--|
| <pre>static std::vector<platform> get_platforms()</platform></pre> | Returns a std::vector containing all |
| | SYCL platforms from all SYCL backends |
| | available in the system. The std::vector |
| | returned must contain at least one SYCL |
| | platform that is from a SYCL host backend. |
| | End of table |

Table 4.11: Static member functions of the SYCL platform class.

4.6.2.2 Platform information descriptors

A platform can be queried for information using the get_info member function of the platform class, specifying one of the info parameters enumerated in info::platform. Every platform (including a host platform) must produce a valid value for each info parameter. The possible values for each info parameter and any restrictions are defined in the specification of the SYCL backend associated with the platform. All info parameters in info::platform are specified in Table 4.12 and the synopsis for info::platform is described in appendix A.1.

| Platform descriptors | Return type | Description |
|---------------------------------------|----------------------------|---|
| <pre>info::platform::version</pre> | <pre>std::string</pre> | Returns the software driver version of the device. |
| <pre>info::platform::name</pre> | <pre>std::string</pre> | Returns the name of the platform. |
| <pre>info::platform::vendor</pre> | <pre>std::string</pre> | Returns the name of the vendor providing the |
| | | platform. |
| <pre>info::platform::extensions</pre> | <pre>std::vector<</pre> | Deprecated, use <pre>device::get_info()</pre> with info |
| | <pre>std::string></pre> | ::device::aspects instead. |
| | | Returns the extensions supported by the platform. |
| | · | End of table |

Table 4.12: Platform information descriptors.

4.6.3 Context class

The context class represents a SYCL context. A context represents the runtime data structures and state required by a SYCL backend API to interact with a group of devices associated with a platform.

The SYCL context class provides the common reference semantics (see Section 4.5.3).

4.6.3.1 Context interface

The constructors and member functions of the SYCL context class are listed in Tables 4.13 and 4.14, respectively. The additional common special member functions and common member functions are listed in 4.5.3 in Tables 4.1 and 4.2, respectively.

All member functions of the context class are synchronous and errors are handled by throwing synchronous SYCL exceptions.

All constructors of the SYCL context class will construct an instance associated with a particular SYCL backend, determined by the constructor parameters or, in the case of the default constructor, the SYCL device produced by the default_selector_v.

A SYCL context can optionally be constructed with an async_handler parameter. In this case the async_handler is used to report asynchronous SYCL exceptions, as described in 4.15.

Information about a SYCL context may be queried through the get_info() member function.

```
1 namespace sycl {
 2
    class context {
 3
     public:
 4
      explicit context(const property_list &propList = {});
 5
 6
      explicit context(async_handler asyncHandler,
 7
                       const property_list &propList = {});
 8
 9
      explicit context(const device &dev, const property_list &propList = {});
10
11
      explicit context(const device &dev, async_handler asyncHandler,
12
                       const property_list &propList = {});
13
14
      explicit context(const std::vector<device> &deviceList,
15
                       const property_list &propList = {});
16
17
      explicit context(const std::vector<device> &deviceList,
18
                       async_handler asyncHandler,
19
                       const property_list &propList = {});
20
      /* -- property interface members -- */
21
22
23
      /* -- common interface members -- */
24
25
      backend get_backend() const;
26
27
      bool is_host() const;
28
29
      platform get_platform() const;
```
| 30 | |
|----|---|
| 31 | <pre>std::vector<device> get_devices() const;</device></pre> |
| 32 | |
| 33 | <pre>template <info::context param=""></info::context></pre> |
| 34 | <pre>typename info::param_traits<info::context, param="">::return_type get_info() const;</info::context,></pre> |
| 35 | |
| 36 | template <typename backendenum="" backendenum,="" param=""></typename> |
| 37 | <pre>typename info::param_traits<backendenum, param="">::return_type</backendenum,></pre> |
| 38 | <pre>get_backend_info() const;</pre> |
| 39 | }; |
| 40 | } // namespace sycl |

| | - |
|--|---|
| Constructor | Description |
| <pre>explicit context(async_handler asyncHandler = {})</pre> | Constructs a SYCL context instance using |
| | an instance of default_selector_v to select |
| | the associated SYCL platform and device |
| | (s). The devices that are associated with the |
| | constructed context are implementation de- |
| | fined but must contain the device chosen by |
| | the device selector. The constructed SYCL |
| | context will use the asyncHandler parame- |
| | ter to handle exceptions. |
| <pre>explicit context(const device &dev,</pre> | Constructs a SYCL context instance using |
| <pre>async_handler asyncHandler = {})</pre> | the dev parameter as the associated SYCL |
| | device and the SYCL platform associated |
| | with the dev parameter as the associated |
| | SYCL platform. The constructed SYCL |
| | context will use the asyncHandler param- |
| | eter to handle exceptions. |
| <pre>explicit context(const std::vector<device> &</device></pre> | Constructs a SYCL context instance us- |
| deviceList, | ing the SYCL device(s) in the deviceList |
| <pre>async_handler asyncHandler = {})</pre> | parameter as the associated SYCL device |
| | (s) and the SYCL platform associated with |
| | each SYCL device in the deviceList pa- |
| | rameter as the associated SYCL platform |
| | . This requires that all SYCL devices in |
| | the deviceList parameter have the same as- |
| | sociated SYCL platform. The constructed |
| | SYCL context will use the asyncHandler |
| | parameter to handle exceptions. |
| | End of table |

| | Table 4.13: | Constructors | of the | SYCL | context | class. |
|--|-------------|--------------|--------|------|---------|--------|
|--|-------------|--------------|--------|------|---------|--------|

| Member function | Description |
|---------------------------------------|---|
| <pre>backend get_backend()const</pre> | Returns a backend identifying the SYCL |
| | backend associated with this context. |
| <pre>bool is_host()const</pre> | Returns true if the backend associated with |
| | this SYCL context is a SYCL host backend. |
| | Continued on next page |

 Table 4.14: Member functions of the context class.

 CHAPTER 4. SYCL PROGRAMMING INTERFACE

| Member function | Description |
|--|---|
| <pre>template <info::context param=""> typename info::</info::context></pre> | Queries this SYCL context for information |
| <pre>param_traits<info::context, param="">::return_type</info::context,></pre> | requested by the template parameter param |
| <pre>get_info()const</pre> | using the param_traits class template to |
| | facilitate returning the appropriate type as- |
| | sociated with the param parameter. |
| template <typename backendenum="" backendenum,="" param=""></typename> | Queries this SYCL context for SYCL back- |
| <pre>typename info::param_traits<backendenum, param="">::</backendenum,></pre> | end-specific information requested by the |
| return_type | template parameter param. BackendEnum |
| <pre>get_backend_info()const</pre> | can be any enum class type specified |
| | by the SYCL backend specification of a |
| | supported SYCL backend named accord- |
| | ing to the convention info:: <backend_name< td=""></backend_name<> |
| | >::context and param must be a valid |
| | enumeration of that enum class. Spe- |
| | cializations of info::param_traits must |
| | be defined for BackendEnum in accor- |
| | dance with the SYCL backend specifica- |
| | tion. Must throw an exception with the |
| | errc::invalid_object_error error code if |
| | the SYCL backend that corresponds with |
| | BackendEnum is different from the SYCL |
| | backend that is associated with this context. |
| <pre>platform get_platform()const</pre> | Returns the SYCL platform that is asso- |
| | ciated with this SYCL context. The value |
| | returned must be equal to that returned by |
| | <pre>get_info<info::context::platform>().</info::context::platform></pre> |
| <pre>std::vector<device></device></pre> | Returns a std::vector containing all SYCL |
| <pre>get_devices()const</pre> | devices that are associated with this SYCL |
| | context. The value returned must be |
| | equal to that returned by get_info <info::< td=""></info::<> |
| | <pre>context::devices>().</pre> |
| | End of table |

Table 4.14: Member functions of the context class.

4.6.3.2 Context information descriptors

A context can be queried for information using the get_info member function of the context class, specifying one of the info parameters enumerated in info::context. Every context (including a host context) must produce a valid value for each info parameter. The possible values for each info parameter and any restrictions are defined in the specification of the SYCL backend associated with the context. All info parameters in info::context are specified in Table 4.15 and the synopsis for info::context is described in appendix A.2.

| Context Descriptors | Return type | Description |
|------------------------------------|---|------------------------------------|
| <pre>info::context::platform</pre> | platform | Returns the platform associated |
| | | with the context. |
| info::context::devices | <pre>std::vector<device></device></pre> | Returns all of the devices associ- |
| | | ated with the context. |
| | | Continued on next page |

| Context Descriptors | Return type | Description |
|--|---|--------------------------------------|
| <pre>info::context::</pre> | <pre>std::vector<memory_order></memory_order></pre> | Returns the set of memory or- |
| atomic_memory_order_capabilities | 5 | derings supported by atomic oper- |
| | | ations on all devices in the con- |
| | | text, which is guaranteed to include |
| | | relaxed. |
| | | The memory ordering of the context |
| | | determines the behavior of atomic |
| | | operations applied to any memory |
| | | that can be concurrently accessed |
| | | by multiple devices in the context. |
| <pre>info::context::</pre> | <pre>std::vector<memory_order></memory_order></pre> | Returns the set of memory order- |
| <pre>atomic_fence_order_capabilities</pre> | | ings supported by atomic_fence on |
| | | all devices in the context, which is |
| | | guaranteed to include relaxed. |
| | | The memory ordering of the context |
| | | determines the behavior of fence |
| | | operations applied to any memory |
| | | that can be concurrently accessed |
| | | by multiple devices in the context. |
| <pre>info::context::</pre> | <pre>std::vector<memory_scope></memory_scope></pre> | Returns the set of memory scopes |
| atomic_memory_scope_capabilities | 5 | supported by atomic operations on |
| | | all devices in the context, which is |
| | | guaranteed to include work_group. |
| <pre>info::context::</pre> | <pre>std::vector<memory_scope></memory_scope></pre> | Returns the set of memory order- |
| <pre>atomic_fence_scope_capabilities</pre> | | ings supported by atomic_fence on |
| | | all devices in the context, which is |
| | | guaranteed to include work_group. |
| | | End of table |

| Table 4.15: Context information descriptors |
|---|
|---|

4.6.3.3 Context properties

The property_list constructor parameters are present for extensibility.

4.6.4 Device class

The SYCL device class encapsulates a single SYCL device on which kernels can be executed. A SYCL device object can map to a native backend object.

All member functions of the device class are synchronous and errors are handled by throwing synchronous SYCL exceptions.

The default constructor of the SYCL device class will construct a host device from the host SYCL backend.

The explicit constructor of the SYCL device class which takes a device selector will construct a device selected by the device selector according to Section 4.6.1.1.

A SYCL device can be partitioned into multiple SYCL devices, by calling the create_sub_devices() member function template. The resulting SYCL devices are considered sub devices, and it is valid to partition these sub

devices further. The range of support for this feature is **SYCL** backend and device specific and can be queried for through get_info().

The SYCL device class provides the common reference semantics (see Section 4.5.3).

4.6.4.1 Device interface

A synopsis of the SYCL device class is provided below. The constructors, member functions and static member functions of the SYCL device class are listed in Tables 4.16, 4.17 and 4.18 respectively. The additional common special member functions and common member functions are listed in 4.5.3 in Tables 4.1 and 4.2, respectively.

```
1
    namespace sycl {
 2
 3
    class device {
 4
     public:
 5
      device();
 6
 7
      template <typename DeviceSelector>
 8
      explicit device(const DeviceSelector &deviceSelector);
 9
10
      /* -- common interface members -- */
11
12
      backend get_backend() const;
13
14
      bool is_host() const;
15
16
      bool is_cpu() const;
17
18
      bool is_gpu() const;
19
20
      bool is_accelerator() const;
21
22
      platform get_platform() const;
23
24
      template <info::device param>
25
      typename info::param_traits<info::device, param>::return_type
26
      get_info() const;
27
28
      template <typename BackendEnum, BackendEnum param>
29
      typename info::param_traits<BackendEnum, param>::return_type
30
      get_backend_info() const;
31
32
      bool has(aspect asp) const;
33
34
      bool has_extension(const std::string &extension) const; // Deprecated
35
36
      // Available only when prop == info::partition_property::partition_equally
37
      template <info::partition_property prop>
38
      std::vector<device> create_sub_devices(size_t nbSubDev) const;
39
40
      // Available only when prop == info::partition_property::partition_by_counts
41
      template <info::partition_property prop>
42
      std::vector<device> create_sub_devices(const std::vector<size_t> &counts) const;
43
44
      // Available only when prop == info::partition_property::partition_by_affinity_domain
```

```
45 template <info::partition_property prop>
```

- 46 std::vector<device> create_sub_devices(info::affinity_domain affinityDomain) const;
- 47 48

```
static std::vector<device> get_devices(
```

49 info::device_type deviceType = info::device_type::all);

50 };

```
51 } // namespace sycl
```

| Constructor | Description |
|---|--|
| device() | Constructs a SYCL device instance as a |
| | host device. |
| <pre>template <typename deviceselector=""></typename></pre> | Constructs a SYCL device instance using |
| <pre>explicit device(const DeviceSelector &</pre> | the device selected by the device selector |
| deviceSelector) | provided. |
| | End of table |

Table 4.16: Constructors of the SYCL device class.

| Member function | Description |
|--|--|
| <pre>backend get_backend()const</pre> | Returns the a backend identifying the SYCL |
| | backend associated with this device. |
| <pre>platform get_platform()const</pre> | Returns the associated SYCL platform |
| | . The value returned must be equal to |
| | <pre>that returned by get_info<info::device::< pre=""></info::device::<></pre> |
| | <pre>platform>().</pre> |
| <pre>bool is_host()const</pre> | Returns the same value as has(aspect:: |
| | host). See Table 4.20. |
| <pre>bool is_cpu()const</pre> | Returns the same value as has (aspect::cpu |
| |). See Table 4.20. |
| <pre>bool is_gpu()const</pre> | Returns the same value as has(aspect::gpu |
| |). See Table 4.20. |
| <pre>bool is_accelerator()const</pre> | Returns the same value as has(aspect:: |
| | accelerator). See Table 4.20. |
| <pre>template <info::device param=""> typename info::</info::device></pre> | Queries this SYCL device for information |
| <pre>param_traits<info::device, param="">::return_type</info::device,></pre> | requested by the template parameter param |
| <pre>get_info()const</pre> | . Specializations of info::param_traits |
| | must be defined in accordance with the info |
| | parameters in Table 4.19 to facilitate return- |
| | ing the type associated with the param pa- |
| | rameter. |
| | Continued on next page |

Table 4.17: Member functions of the SYCL device class.

| Member function | Description |
|---|--|
| <pre>template <typename backendenum="" backendenum,="" param=""></typename></pre> | Queries this SYCL device for SYCL back- |
| <pre>typename info::param_traits<backendenum, param="">::</backendenum,></pre> | end-specific information requested by the |
| return_type | template parameter param. BackendEnum |
| get_backend_info()const | can be any enum class type specified |
| | by the SYCL backend specification of a |
| | supported SYCL backend named accord- |
| | ing to the convention info:: <backend_name< th=""></backend_name<> |
| | >::device and param must be a valid |
| | enumeration of that enum class. Spe- |
| | cializations of info::param_traits must |
| | be defined for BackendEnum in accor- |
| | dance with the SYCL backend specifica- |
| | tion. Must throw an exception with the |
| | errc::invalid_object_error error code if |
| | the SYCL backend that corresponds with |
| | BackendEnum is different from the SYCL |
| | backend that is associated with this device |
| hool has(aspect asp)const | Returns true if this SYCL device has the |
| boor has (aspect asp) const | given aspect SYCL applications can use |
| | this member function to determine which |
| | optional features this device supports (if |
| | any). |
| <pre>bool has_extension (const std::string &extension)</pre> | Deprecated, use has () instead. |
| const | Returns true if this SYCL device supports |
| | the extension queried by the extension pa- |
| | rameter. |
| <pre>template <info::partition_property prop=""></info::partition_property></pre> | Available only when prop is info:: |
| <pre>std::vector<device> create_sub_devices(</device></pre> | <pre>partition_property::partition_equally</pre> |
| <pre>size_t nbSubDev)const</pre> | . Returns a std::vector of sub devices |
| | partitioned from this SYCL device |
| | equally based on the nbSubDev param- |
| | eter. If this SYCL device does not |
| | <pre>support info::partition_property::</pre> |
| | partition_equally an exception with the |
| | errc::feature_not_supported error code |
| | must be thrown. |
| template <info::partition_property prop=""></info::partition_property> | Available only when prop is |
| const std::vector <size t=""> &counts)const</size> | nartition by count Returns a std: |
| | vector of sub devices partitioned from this |
| | SYCL device by count sizes based on the |
| | counts parameter. If the SYCL device does |
| | not support info::partition_property:: |
| | partition_by_count an exception with the |
| | errc::feature_not_supported error code |
| | must be thrown. |
| | Continued on next page |

Table 4.17: Member functions of the SYCL device class.

| Member function | Description |
|---|---|
| <pre>template <info::partition_property prop=""></info::partition_property></pre> | Available only when prop is |
| <pre>std::vector<device> create_sub_devices(</device></pre> | <pre>info::partition_property::</pre> |
| <pre>info::affinity_domain affinityDomain)const</pre> | partition_by_affinity_domain. Returns |
| | a std::vector of sub devices parti- |
| | tioned from this SYCL device by affinity |
| | domain based on the affinityDomain |
| | parameter. Partitions the device into |
| | sub devices based upon the affinity do- |
| | main. If the SYCL device does not |
| | <pre>support info::partition_property::</pre> |
| | partition_by_affinity_domain or the |
| | SYCL device does not support the info:: |
| | affinity_domain provided, an exception |
| | with the errc::feature_not_supported |
| | error code must be thrown. |
| | End of table |

Table 4.17: Member functions of the SYCL device class.

| Static member function | Description |
|--|---|
| <pre>static std::vector<device></device></pre> | Returns a std::vector containing all SYCL |
| get_devices(| devices from all SYCL backends available |
| <pre>info::device_type deviceType =</pre> | in the system of the device type specified by |
| <pre>info::device_type::all)</pre> | the parameter deviceType. Note that when |
| | <pre>the device_type is info::device_type::</pre> |
| | all or info::device_type::host, the std |
| | ::vector returned must contain at least one |
| | host device from the SYCL host backend. |
| | End of table |

Table 4.18: Static member functions of the SYCL device class.

4.6.4.2 Device information descriptors

A device can be queried for information using the get_info member function of the device class, specifying one of the info parameters enumerated in info::device. Every device (including a host device) must produce a valid value for each info parameter. The possible values for each info parameter and any restriction are defined in the specification of the SYCL backend associated with the device. All info parameters in info::device are specified in Table 4.19 and the synopsis for info::device is described in appendix A.3.

| Device descriptors | Return type | Description |
|--------------------------------------|-------------|---|
| <pre>info::device::device_type</pre> | info:: | Returns the device type associated with the |
| | device_type | <pre>device. May not return info::device_type::</pre> |
| | | all. |
| <pre>info::device::vendor_id</pre> | uint32_t | Returns a unique vendor device identifier. |
| | | Continued on next page |

Table 4.19: Device information descriptors.

| Device descriptors | Return type | Description |
|--|--------------|---|
| <pre>info::device::</pre> | uint32_t | Returns the number of parallel compute units |
| <pre>max_compute_units</pre> | | available to the device. The minimum value is |
| | | 1. |
| <pre>info::device::</pre> | uint32_t | Returns the maximum dimensions that specify |
| <pre>max_work_item_dimensions</pre> | | the global and local work-item IDs used by the |
| | | data parallel execution model. The minimum |
| | | value is 3 if this SYCL device is not of device |
| | | <pre>type info::device_type::custom.</pre> |
| <pre>info::device::</pre> | id<3> | Returns the maximum number of work-items |
| <pre>max_work_item_sizes</pre> | | that are permitted in each dimension of the |
| | | work-group of the nd_range. The minimum |
| | | value is (1, 1, 1) for devices that are not of de- |
| | | vice type info::device_type::custom. |
| info::device:: | size_t | Returns the maximum number of work-items |
| max_work_group_size | | that are permitted in a work-group executing a |
| | | kernel on a single compute unit. The minimum |
| | 1 | Value is 1. |
| info::device:: | uint32_t | Returns the maximum number of sub-groups |
| max_num_sub_groups | | daviae. The minimum value is 1 |
| infodoui.co | heel | Deturns true if the device supports indepen |
| sub group independent forward p | DOOL | dent forward progress of sub groups with re- |
| Sub_group_independent_iorward_p | LOGIESS | spect to other sub-groups in the same work- |
| | | group |
| info::device::sub group sizes | std::vector< | Returns a std::vector of size t containing |
| into i acvice i bab_gi bap_bileb | size t> | the set of sub-group sizes supported by the de- |
| | 011020 | vice. |
| info::device:: | uint32 t | Returns the preferred native vector width |
| preferred_vector_width_char | | size for built-in scalar types that can be put |
| info::device:: | | into vectors. The vector width is defined as |
| preferred_vector_width_short | | the number of scalar elements that can be |
| <pre>info::device::</pre> | | stored in the vector. Must return 0 for info:: |
| preferred_vector_width_int | | <pre>device::preferred_vector_width_double if</pre> |
| <pre>info::device::</pre> | | the device does not have aspect::fp64 |
| <pre>preferred_vector_width_long</pre> | | and must return 0 for info::device:: |
| <pre>info::device::</pre> | | preferred_vector_width_half if the device |
| <pre>preferred_vector_width_float</pre> | | does not have aspect::fp16. |
| info::device:: | | |
| <pre>preferred_vector_width_double</pre> | | |
| info::device:: | | |
| preferred_vector_width_half | | |
| | | Continued on next page |

| Device descriptors | Return type | Description |
|--|-------------|---|
| info::device:: | uint32_t | Returns the native ISA vector width. |
| native_vector_width_char | | The vector width is defined as the num- |
| <pre>info::device::</pre> | | ber of scalar elements that can be stored |
| native_vector_width_short | | in the vector. Must return 0 for info:: |
| <pre>info::device::</pre> | | <pre>device::preferred_vector_width_double if</pre> |
| <pre>native_vector_width_int</pre> | | the device does not have aspect::fp64 |
| <pre>info::device::</pre> | | and must return 0 for info::device:: |
| <pre>native_vector_width_long</pre> | | preferred_vector_width_half if the device |
| <pre>info::device::</pre> | | does not have aspect::fp16. |
| <pre>native_vector_width_float</pre> | | |
| <pre>info::device::</pre> | | |
| <pre>native_vector_width_double</pre> | | |
| <pre>info::device::</pre> | | |
| native_vector_width_half | | |
| info::device:: | uint32_t | Returns the maximum configured clock fre- |
| <pre>max_clock_frequency</pre> | | quency of this SYCL device in MHz. |
| <pre>info::device::address_bits</pre> | uint32_t | Returns the default compute device address |
| | | space size specified as an unsigned integer |
| | | value in bits. Must return either 32 or 64. |
| info::device:: | uint64_t | Returns the maximum size of memory object |
| <pre>max_mem_alloc_size</pre> | | allocation in bytes. The minimum value is max |
| | | (1/4th of info::device::global_mem_size |
| | | ,128*1024*1024) if this SYCL device is not |
| | | of device type info::device_type::custom. |
| <pre>info::device::image_support</pre> | bool | Deprecated. |
| | | Returns the same value as device::has(|
| | | aspect::image). |
| info::device:: | uint32_t | Returns the maximum number of simultane- |
| <pre>max_read_image_args</pre> | | ous image objects that can be read from by |
| | | a kernel. The minimum value is 128 if the |
| | | SYCL device has aspect::image. |
| info::device:: | uint32_t | Returns the maximum number of simultane- |
| <pre>max_write_image_args</pre> | | ous image objects that can be written to by a |
| | | kernel. The minimum value is 8 if the SYCL |
| | | device has aspect::image. |
| info::device:: | size_t | Returns the maximum width of a 2D image |
| image2d_max_width | | or 1D image in pixels. The minimum value is |
| | | 8192 if the SYCL device has aspect::image. |
| info::device:: | size_t | Returns the maximum height of a 2D image |
| <pre>image2d_max_height</pre> | | in pixels. The minimum value is 8192 if the |
| | | SYCL device has aspect::image. |
| info::device:: | size_t | Returns the maximum width of a 3D image |
| image3d_max_width | | in pixels. The minimum value is 2048 if the |
| | | SYCL device has aspect::image. |
| info::device:: | size_t | Returns the maximum height of a 3D image |
| <pre>image3d_max_height</pre> | | in pixels. The minimum value is 2048 if the |
| | | SYCL device has aspect::image. |
| | | Continued on next page |

| Device descriptors | Return type | Description |
|---------------------------------------|-------------|--|
| info::device:: | size_t | Returns the maximum depth of a 3D image |
| image3d_max_depth | | in pixels. The minimum value is 2048 if the |
| | | SYCL device has aspect::image. |
| info::device:: | size_t | Returns the number of pixels for a 1D im- |
| <pre>image_max_buffer_size</pre> | | age created from a buffer object. The mini- |
| | | mum value is 65536 if the SYCL device has |
| | | aspect::image. Note that this information is |
| | | intended for OpenCL interoperability only as |
| | | this feature is not supported in SYCL. |
| <pre>info::device::</pre> | size_t | Returns the maximum number of images in a |
| <pre>image_max_array_size</pre> | | 1D or 2D image array. The minimum value is |
| | | 2048 if the SYCL device has aspect::image. |
| <pre>info::device::max_samplers</pre> | uint32_t | Returns the maximum number of samplers |
| | | that can be used in a kernel. The minimum |
| | | value is 16 if the SYCL device has aspect:: |
| | | image. |
| <pre>info::device::</pre> | size_t | Returns the maximum size in bytes of the ar- |
| <pre>max_parameter_size</pre> | | guments that can be passed to a kernel. The |
| | | minimum value is 1024 if this SYCL device |
| | | is not of device type info::device_type:: |
| | | custom. For this minimum value, only a maxi- |
| | | mum of 128 arguments can be passed to a ker- |
| | | nel. |
| <pre>info::device::</pre> | uint32_t | Returns the minimum value in bits of the |
| <pre>mem_base_addr_align</pre> | | largest supported SYCL built-in data type if |
| | | this SYCL device is not of device type info |
| | | ::device_type::custom. |
| | | Continued on next page |

Table 4.19: Device information descriptors.

| Device descriptors | Return type | Description |
|---|--|--|
| <pre>info::device::half_fp_config</pre> | <pre>std::vector< info::fp_config</pre> | Returns a std::vector of info::fp_config describing the half precision floating-point |
| | > | capability of this SYCL device. The std:: |
| | | vector may contain zero or more of the fol- |
| | | lowing values: |
| | | info::fp_config::denorm: denorms are supported. |
| | | info::fp_config::inf_nan: INF and quiet NaNs are supported. |
| | | info::fp_config::round_to_nearest: round to nearest even rounding mode is supported. |
| | | info::fp_config::round_to_zero round to zero rounding mode is supported. |
| | | info::fp_config::round_to_inf: round to positive and negative infinity rounding modes are supported. |
| | | info::fp_config::fma: IEEE754- 2008 fused multiply add is supported. |
| | | info::fp config:: |
| | | correctly_rounded_divide_sqrt: |
| | | divide and sqrt are correctly rounded as |
| | | defined by the IEEE754 specification. |
| | | <pre>• info::fp_config::soft_float: basic</pre> |
| | | floating-point operations (such as addi- |
| | | tion, subtraction, multiplication) are im- |
| | | plemented in software. |
| | | If half precision is supported by this SYCL |
| | | device (i.e. the device has aspect::fp16) |
| | | there is no minimum floating-point capability. |
| | | If half support is not supported the returned |
| | | std::vector must be empty. |
| | | Continued on next page |

| Device descriptors | Return type | Description |
|--|---|---|
| Device descriptors info::device::single_fp_config | <pre>Return type std::vector< info::fp_config ></pre> | <pre>Description Returns a std::vector of info::fp_config describing the single precision floating-point capability of this SYCL device. The std:: vector must contain one or more of the fol- lowing values: info::fp_config::denorm: denorms are supported. info::fp_config::inf_nan: INF and quiet NaNs are supported. info::fp_config::round_to_nearest: round to nearest even rounding mode is supported. info::fp_config::round_to_zero : round to zero rounding mode is supported. info::fp_config::round_to_inf: round to zero rounding mode is supported. info::fp_config::round_to_inf: round to positive and negative infinity rounding modes are supported. info::fp_config::fma: IEEE754- 2008 fused multiply add is supported. info::fp_config:: correctly_rounded_divide_sqrt: divide and sqrt are correctly rounded as defined by the IEEE754 specification. info::fp_config::soft_float: basic floating-point operations (such as addi- tion, subtraction, multiplication) are im- plemented in software. If this SYCL device is not of type info ::device_type::custom then the minimum floating-point capability must be:info:: </pre> |
| | | floating-point capability must be: info:: |
| | | <pre>tp_config::round_to_nearest and info:: fm_config::inf_non</pre> |
| | | <pre>tp_contig::int_nan.</pre> |
| | | Continued on next page |

| Device descriptors | Return type | Description |
|--|----------------------------|--|
| <pre>info::device::double_fp_config</pre> | <pre>std::vector<</pre> | Returns a std::vector of info::fp_config |
| | info::fp_config | describing the double precision floating-point |
| | > | capability of this SYCL device. The std:: |
| | | vector may contain zero or more of the fol- |
| | | lowing values: |
| | | • info::fp_config::denorm: denorms |
| | | are supported. |
| | | • info:::tp_config:::nf_nan: INF and |
| | | in factor for the supported. |
| | | Info:::p_config::round_to_nearest. round to nearest even rounding mode is |
| | | supported |
| | | • info::fn config::round to zero |
| | | : round to zero rounding mode is |
| | | supported. |
| | | <pre>• info::fp_config::round_to_inf:</pre> |
| | | round to positive and negative infinity |
| | | rounding modes are supported. |
| | | • info::fp_config::fma: IEEE754- |
| | | 2008 fused multiply-add is supported. |
| | | info::fp_config::soft_float: basic |
| | | floating-point operations (such as addi- |
| | | tion, subtraction, multiplication) are im- |
| | | plemented in software. |
| | | If double precision is supported by this SVCL doubles (i.e. the doubles has accest |
| | | sich device (i.e. the device has aspect |
| | | type info: device type: custom then the |
| | | minimum floating-point capability must be: |
| | | info::fp config::fma, info::fp config |
| | | ::round to nearest. info::fp config |
| | | ::round_to_zero, info::fp_config:: |
| | | round_to_inf, info::fp_config::inf_nan |
| | | and info::fp_config::denorm. If dou- |
| | | ble support is not supported the returned |
| | | std::vector must be empty. |
| <pre>info::device::</pre> | info:: | Returns the type of global memory cache sup- |
| global_mem_cache_type | global_mem_cache | ported. |
| | type | |
| <pre>info::device:: .label mem ersby line sint</pre> | uint32_t | Returns the size of global memory cache line |
| giobal_mem_cache_line_size | uin+64 + | III Uyits. Returns the size of global memory cache in |
| alobal mem cache size | u11104_C | hytes |
| info::device::global mem size | uint64 t | Returns the size of global device memory in |
| | | bytes. |
| | 1 | Continued on next page |

| int64_t | Returns the maximum size in bytes of a con- |
|---------------|---|
| | Returns the maximum size in bytes of a con- |
| | stant buffer allocation. The minimum value is |
| | 64 KB if this SYCL device is not of type info |
| | ::device_type::custom. |
| int32_t | Returns the maximum number of constant ar- |
| | guments that can be declared in a kernel. The |
| | minimum value is 8 if this SYCL device is not |
| | of type info::device_type::custom. |
| nfo:: | Returns the type of local memory supported. |
| ocal_mem_type | This can be info::local_mem_type::local |
| | implying dedicated local memory storage |
| | such as SRAM, or info::local_mem_type:: |
| | global. If this SYCL device is of type |
| | info::device_type::custom this can also be |
| | info::local_mem_type::none, indicating lo- |
| | Paturns, the size of level memory arene. |
| 111104_1 | in bytes. The minimum value is 32 KB |
| | if this SVCL device is not of type info: |
| | device type::custom |
| 00] | Returns true if the device implements error |
| | correction for all accesses to compute device |
| | memory (global and constant). Returns false if |
| | the device does not implement such error cor- |
| | rection. |
| ool | Deprecated, use device::has() with one of |
| | the aspect::usm_* aspects instead. |
| | Returns true if the device and the host have |
| | a unified memory subsystem and returns false |
| | otherwise. |
| td::vector< | Returns the set of memory orderings sup- |
| emory_order> | ported by atomic operations on the device, |
| | which is guaranteed to include relaxed. |
| | If this device is a host device, the set must |
| | include all values of the memory_order enum |
| | class: relaxed, acquire, release, acq_rel |
| • | and seq_cst. |
| tu::vector< | Returns the set of memory orderings sup- |
| emory_order> | is guaranteed to include released |
| | Is guardineed to include relaxed. |
| | include all values of the momenty order crum |
| | class: relayed acquire relass acc rel |
| | and sea est |
| | Continued on next page |
| | nt32_t fo:: cal_mem_type nt64_t nol cd::vector< emory_order> cd::vector< emory_order> |

| Device descriptors | Return type | Description |
|--|--|---|
| info::device:: | <pre>std::vector<</pre> | Returns the set of memory scopes supported |
| atomic_memory_scope_capabilities | s memory_scope> | by atomic operations on the device, which is |
| | | guaranteed to include work_group. |
| | | If this device is a host device, the set must |
| | | include all values of the memory_scope enum |
| | | <pre>class: work_item, sub_group, work_group,</pre> |
| | | device and system. |
| <pre>info::device::</pre> | <pre>std::vector<</pre> | Returns the set of memory scopes supported |
| <pre>atomic_fence_scope_capabilities</pre> | <pre>memory_scope></pre> | by atomic_fence on the device, which is guar- |
| | | anteed to include work_group. |
| | | If this device is a host device, the set must |
| | | include all values of the memory_scope enum |
| | | <pre>class: work_item, sub_group, work_group,</pre> |
| | | device and system. |
| <pre>info::device::</pre> | size_t | Returns the resolution of device timer in |
| profiling_timer_resolution | | nanoseconds. |
| <pre>info::device::is_endian_little</pre> | bool | Returns true if this SYCL device is a little |
| | | endian device and returns false otherwise. |
| <pre>info::device::is_available</pre> | bool | Returns true if the SYCL device is available |
| | | and returns false if the device is not available. |
| <pre>info::device::</pre> | bool | Deprecated. |
| is_compiler_available | | Returns the same value as device::has(|
| | | <pre>aspect::online_compiler).</pre> |
| <pre>info::device::</pre> | bool | Deprecated. |
| is_linker_available | | Returns the same value as device::has(|
| | | <pre>aspect::online_linker).</pre> |
| <pre>info::device::</pre> | <pre>std::vector</pre> | Returns a std::vector of the info:: |
| execution_capabilities | <info::< td=""><td>execution_capability describing the sup-</td></info::<> | execution_capability describing the sup- |
| | execution | ported execution capabilities. Note that this |
| | capability> | information is intended for OpenCL interop- |
| | | erability only as SYCL only supports info:: |
| | | <pre>execution_capability::exec_kernel.</pre> |
| <pre>info::device::queue_profiling</pre> | bool | Deprecated. |
| | | Returns the same value as device::has(|
| | | <pre>aspect::queue_profiling).</pre> |
| <pre>info::device::built_in_kernels</pre> | <pre>std::vector<std< pre=""></std<></pre> | Returns a std::vector of built-in OpenCL ker- |
| | ::string> | nels supported by this SYCL device. |
| <pre>info::device::platform</pre> | platform | Returns the SYCL platform associated with |
| | | this SYCL device. |
| <pre>info::device::name</pre> | <pre>std::string</pre> | Returns the device name of this SYCL device. |
| <pre>info::device::vendor</pre> | <pre>std::string</pre> | Returns the vendor of this SYCL device. |
| <pre>info::device::driver_version</pre> | <pre>std::string</pre> | Returns the OpenCL software driver version |
| | | as a std::string in the form: major_num- |
| | | ber.minor_number, if this SYCL device is an |
| | | OpenCL device. Must return a std::string |
| | | with the value "1.2" if this SYCL device is a |
| | | host device. |

Continued on next page

Table 4.19: Device information descriptors.

| Device descriptors | Return type | Description |
|--|----------------------------|--|
| <pre>info::device::profile</pre> | <pre>std::string</pre> | Returns the OpenCL profile as a std::string |
| | | , if this SYCL device is an OpenCL device. |
| | | The value returned can be one of the following |
| | | strings: |
| | | • FULL_PROFILE - if the device supports |
| | | the OpenCL specification (functionality |
| | | defined as part of the core specification |
| | | and does not require any extensions to |
| | | be supported). |
| | | • EMBEDDED_PROFILE - if the device |
| | | supports the OpenCL embedded profile. |
| | | Must return a std::string with the value |
| | | "FULL PROFILE" if this is a host device. |
| <pre>info::device::version</pre> | <pre>std::string</pre> | Returns the SYCL version as a std:: |
| | | string in the form: <major_version>.<</major_version> |
| | | minor_version>. If this SYCL device |
| | | is a host device, the <major_version>.<</major_version> |
| | | minor_version> value returned must be "1.2 |
| | | ". |
| <pre>info::device::backend_version</pre> | <pre>std::string</pre> | Returns a string describing the version of |
| | | the SYCL backend associated with the device. |
| | | The possible values are specified in the SYCL |
| | | backend specification of the SYCL backend |
| | | associated with the device. |
| <pre>info::device::aspects</pre> | <pre>std::vector<</pre> | Returns a std::vector of aspect values sup- |
| | aspect> | ported by this SYCL device. |
| | | Continued on next page |

| Device descriptors | Return type | Description |
|-------------------------------------|--|---|
| <pre>info::device::extensions</pre> | <pre>std::vector<std< pre=""></std<></pre> | Deprecated, use info::device::aspects in- |
| | ::string> | stead. |
| | | Returns a std::vector of extension names |
| | | (the extension names do not contain any |
| | | spaces) supported by this SYCL device. The |
| | | extension names returned can be vendor sup- |
| | | ported extension names and one or more of the |
| | | following Khronos approved extension names: |
| | | cl_khr_int64_base_atomics |
| | | cl_khr_int64_extended_atomics |
| | | cl_khr_3d_image_writes cl_khr_5d_1 |
| | | • Cl_Knr_Iplo |
| | | • Cl_Knr_gl_Sharing |
| | | • cl_knr_g1_event |
| | | • cl_khr_dv0_media_sharing |
| | | • cl_khr_d3d11_sharing |
| | | • cl khr denth images |
| | | • cl khr gl depth images |
| | | • cl khr gl msaa sharing |
| | | cl_khr_image2d_from_buffer |
| | | cl_khr_initialize_memory |
| | | cl_khr_context_abort |
| | | cl_khr_spir |
| | | If this SYCL device is an OpenCL device then |
| | | following approved Khronos extension names |
| | | must be returned by all device that support |
| | | OpenCL C 1.2: |
| | | cl_khr_global_int32_base_atomics |
| | | cl_khr_global_int32_extended_atomics |
| | | cl_khr_local_int32 base atomics |
| | | • cl_khr_local_int32_extended_atomics |
| | | cl_khr_byte_addressable_store |
| | | • cl_khr_fp64 (for backward compatibil- |
| | | ity if double precision is supported) |
| | | Please refer to the OpenCL 1.2 Extension |
| | | Specification for a detailed description of these |
| | | extensions. |
| <pre>info::device::</pre> | size_t | Returns the maximum size of the internal |
| <pre>printf_buffer_size</pre> | | buffer that holds the output of printf calls |
| | | from a kernel. The minimum value is 1 MB |
| | | if info::device::profile returns true for this |
| | | SYCL device. |
| | | Continued on next page |

| Device descriptors | Return type | Description |
|--|--|--|
| info::device:: | bool | Returns true if the preference for this SYCL |
| <pre>preferred_interop_user_sync</pre> | | device is for the user to be responsible for |
| | | synchronization, when sharing memory ob- |
| | | jects between OpenCL and other APIs such |
| | | as DirectX, false if the device/implementation |
| | | has a performant path for performing synchro- |
| | | nization of memory object shared between |
| | | OpenCL and other APIs such as DirectX. |
| <pre>info::device::parent_device</pre> | device | Returns the parent SYCL device to which |
| | | this sub-device is a child if this is a sub- |
| | | device. Must throw an exception with the |
| | | <pre>errc::invalid_object_error error code if</pre> |
| | | this SYCL device is not a sub device. |
| info::device:: | uint32_t | Returns the maximum number of sub-devices |
| <pre>partition_max_sub_devices</pre> | | that can be created when this SYCL device |
| | | is partitioned. The value returned cannot ex- |
| | | ceed the value returned by info::device:: |
| | | <pre>device_max_compute_units.</pre> |
| <pre>info::device::</pre> | <pre>std::vector</pre> | Returns the partition properties supported |
| partition_properties | <info::< td=""><td>by this SYCL device; a vector of info::</td></info::<> | by this SYCL device; a vector of info:: |
| | partition_prop- | partition_property. If this SYCL device |
| | erty> | cannot be partitioned into at least two sub de- |
| | | vices then the returned vector must be empty. |
| <pre>info::device::</pre> | <pre>std::vector</pre> | Returns a std::vector of the partition affinity |
| partition_affinity_domains | <info::< td=""><td>domains supported by this SYCL device when</td></info::<> | domains supported by this SYCL device when |
| | parition_affini | partitioning with info::partition_property |
| | -ty_domain> | ::parition_by_affinity_domain. |
| <pre>info::device::</pre> | info:: | Returns the partition property of this SYCL |
| partition_type_property | partition_prop- | device. If this SYCL device is not a sub de- |
| | erty | vice then the the return value must be info |
| | | ::partition_property::no_partition, oth- |
| | | erwise it must be one of the following values: |
| | | info::partition_property::partition |
| | | equally |
| | | info::partition_property::partition_by |
| | | counts |
| | | info::partition_property::partition_by |
| | | affinity_domain |
| | | |
| | | Continued on next page |

| Device descriptors | Return type | Description |
|--|---|---|
| Device descriptors info::device:: partition_type_affinity_domain | Return type info:: partition_affi- nity_domain | Description Returns the partition affinity domain of this SYCL device. If this SYCL device is not a sub device or the sub device was not partitioned with info::partition_type ::partition_by_affinity_domain then the the return value must be info::partition_affinity_domain:: not_applicable, otherwise it must be one of the following values: • info::partition_affinity_domain::L4 cache • info::partition_affinity_domain::L3 cache • info::partition_affinity_domain::L2 cache • info::partition_affinity_domain::L2 cache • info::partition_affinity_domain::L1 cache • info::partition_affinity_domain::L1 |
| | | partitionable |
| <pre>info::device::reference_count</pre> | uint32_t | Returns the device reference count. If the device is not a sub-device the value returned must be 1. |
| | | End of table |

Table 4.19: Device information descriptors.

4.6.4.3 Device aspects

Every SYCL device has an associated set of "aspects" which identify characteristics of the device. Aspects are defined via the enum class aspect enumeration:

```
1 namespace sycl {
2
3 enum class aspect {
4 host,
```

- 5 cpu,
- 6 gpu,
- 7 accelerator,
- 8 custom,
- 9 fp16,
- 10 fp64,
- 11 int64_base_atomics,
- 12 int64_extended_atomics,
- 13 image,
- 14 online_compiler,
- 15 online_linker,
- 16 queue_profiling,
- 17 usm_device_allocations,
- 18 usm_host_allocations,

```
19 usm_shared_allocations,
20 usm_restricted_shared_allocations,
21 usm_system_allocator
22 };
23
24 } // namespace sycl
```

Table 4.20 lists the aspects that are defined in the core SYCL specification. However, a SYCL backend or extension may provide additional aspects. If so, the SYCL backend specification document or the extension document describes them. If a SYCL backend defined by the Khronos SYCL group provides aspects, their enumerated values are defined in the backend's namespace. For example, an aspect specific to the OpenCL backend could be defined like this:

```
1 namespace sycl {
2 namespace opencl {
3 namespace aspect {
4
5 static constexpr auto bar = static_cast<sycl::aspect>(-1);
6
7 } // namespace aspect
8 } // namespace opencl
9 } // namespace sycl
```

Aspects provided by an extension or a vendor's SYCL backend are defined as described in Chapter 6.

SYCL applications can query the aspects for a **device** via **device**::has() in order to determine whether the **device** supports any optional features. Table 4.20 tells which optional features are enabled by aspects in the core SYCL specification, but backends and extensions may provide optional features also. If so, the SYCL backend specification document or the extension document describes which features are enabled by each aspect.

A SYCL application can also use the is_aspect_active<aspect>::value trait to test whether an aspect is "active" at compile time. An aspect is active if the compilation environment supports any device which has that aspect. For example, if the implementation supports no devices with aspect::custom, the trait is_aspect_active<aspect::custom>::value will be false. The set of active aspects could also be affected by command line options passed to the compiler. For example, if an implementation provides a command line option that disables aspect:: accelerator devices, that trait will be false when the option is passed to the compiler.

[Note: Like any type trait, the value of is_aspect_active<aspect>::value has a uniform value across all parts of a SYCL application. If an implementation uses SMCP, all compiler passes define a particular aspect's is_aspect_active type trait with the same value, regardless of whether that compiler pass's device supports the aspect. Thus, is_aspect_active cannot be used to determine whether any particular device supports an aspect. Instead, applications must use device::has() or platform::has() for this. — end note]

The trait sycl::is_aspect_active<aspect>::value must be defined as either true or false for all the core SYCL aspects listed in Table 4.20, all aspects from Khronos ratified extensions, and all of a vendor's own extension aspects.

| Aspect | Description |
|---|--|
| aspect::host | A device that runs on the host CPU and |
| | exposes the host SYCL backend. Devices |
| | with this aspect have device type info:: |
| | <pre>device_type::host.</pre> |
| aspect::cpu | A device that runs on a CPU, but doesn't |
| | use the host SYCL backend. Devices |
| | with this aspect have device type info:: |
| | device_type::cpu. |
| aspect::gpu | A device that can also be used to accelerate |
| | a 3D graphics API. Devices with this aspect |
| | have device type info::device_type::gpu. |
| aspect::accelerator | A dedicated accelerator device, usually us- |
| | ing a peripheral interconnect for communi- |
| | cation. Devices with this aspect have device |
| | <pre>type info::device_type::accelerator.</pre> |
| aspect::custom | A dedicated accelerator that can use the |
| | SYCL API, but programmable kernels can- |
| | not be dispatched to the device, only fixed |
| | functionality is available. See Section 3.8.7. |
| | Devices with this aspect have device type |
| | <pre>info::device_type::custom.</pre> |
| aspect::fp16 | Indicates that the device supports half pre- |
| | cision floating point operations. |
| aspect::fp64 | Indicates that the device supports 64-bit pre- |
| | cision floating point operations. |
| <pre>aspect::int64_base_atomics</pre> | Indicates that the device supports the |
| | following atomic operations on 64- |
| | bit values: atomic::load, atomic:: |
| | <pre>store, atomic::fetch_add, atomic::</pre> |
| | <pre>fetch_sub, atomic::exchange, atomic::</pre> |
| | compare_exchange_strong and atomic:: |
| | compare_exchange_weak. |
| <pre>aspect::int64_extended_atomics</pre> | Indicates that the device supports the fol- |
| | lowing atomic operations on 64-bit values: |
| | atomic::fetch_min, atomic::fetch_max, |
| | atomic::fetch_and, atomic::fetch_or, |
| | and atomic::fetch_xor. |
| aspect::image | Indicates that the device supports images |
| | (Section 4.7.3). Devices of type info:: |
| | device_type::nost always nave this sup- |
| | Poll. |
| aspect::online_complier | line compilation of device supports on- |
| | that have this aspect support the ball of |
| | and commite() functions defined in Sec. |
| | tion 4 12 7 |
| | U0II 4.13./. |
| | Continued on next page |

Table 4.20: Device aspects defined by the core SYCL specification.

| Aspect | Description |
|--|--|
| <pre>aspect::online_linker</pre> | Indicates that the device supports online |
| | linking of device code. Devices that have |
| | this aspect support the link() functions |
| | defined in Section 4.13.7. All devices |
| | that have this aspect also have aspect:: |
| | online_compiler. |
| <pre>aspect::queue_profiling</pre> | Indicates that the device supports |
| | queue profiling via property::queue:: |
| | enable_profiling. |
| <pre>aspect::usm_device_allocations</pre> | Indicates that the device supports explicit |
| | USM allocations as described in Section 4.8. |
| <pre>aspect::usm_host_allocations</pre> | Indicates that the device can access USM |
| | memory allocated via usm::alloc::host. |
| | (See Section 4.8.) |
| <pre>aspect::usm_shared_allocations</pre> | Indicates that the device supports USM |
| | memory allocated via usm::alloc::shared |
| | as restricted USM, concurrent USM, or |
| | both. (See Section 4.8.) |
| <pre>aspect::usm_restricted_shared_allocations</pre> | Indicates that the device supports USM |
| | memory allocated via usm::alloc:: |
| | shared as restricted USM. Any device |
| | with this aspect will also have aspect |
| | ::usm_shared_allocations. (See Sec- |
| | tion 4.8.) |
| <pre>aspect::usm_system_allocator</pre> | Indicates that the system allocator may |
| | be used instead of SYCL USM allocation |
| | mechanisms for usm::alloc::shared allo- |
| | cations on this device. (See Section 4.8.) |
| | End of table |

Table 4.20: Device aspects defined by the core SYCL specification.

4.6.5 Queue class

The SYCL queue class encapsulates a single SYCL queue which schedules kernels on a SYCL device. The SYCL queue can encapsulate to one or multiple native backend objects.

A SYCL queue can be used to submit command groups to be executed by the SYCL runtime using the submit member function.

All member functions of the queue class are synchronous and errors are handled by throwing synchronous SYCL exceptions. The submit member function schedules command groups asynchronously, so any errors in the submission of a command group are handled by throwing synchronous SYCL exceptions. Any exceptions from the command group after it has been submitted are handled by passing asynchronous errors at specific times to an async_handler, as described in 4.15.

A SYCL queue can wait for all command groups that it has submitted by calling wait or wait_and_throw.

The default constructor of the SYCL queue class will construct a queue based on the SYCL device returned from the default_selector_v (see Section 4.6.1.1).

CHAPTER 4. SYCL PROGRAMMING INTERFACE

94

All other constructors construct a host or device queue, determined by the parameters provided. All constructors will implicitly construct a SYCL platform, device and context in order to facilitate the construction of the queue.

Each constructor takes as the last parameter an optional SYCL property_list to provide properties to the SYCL queue.

The SYCL queue class provides the common reference semantics (see Section 4.5.3).

4.6.5.1 Queue interface

A synopsis of the SYCL queue class is provided below. The constructors and member functions of the SYCL queue class are listed in Tables 4.21 and 4.22 respectively. The additional common special member functions and common member functions are listed in 4.5.3 in Tables 4.1 and 4.2, respectively.

```
namespace sycl {
 1
2
    class queue {
3
     public:
 4
      explicit queue(const property_list &propList = {});
 5
 6
      explicit queue(const async_handler &asyncHandler,
7
                     const property_list &propList = {});
8
9
      template <typename DeviceSelector>
10
      explicit queue(const DeviceSelector &deviceSelector,
11
                     const property_list &propList = {});
12
13
      template <typename DeviceSelector>
14
      explicit queue(const DeviceSelector &deviceSelector,
15
                     const async_handler &asyncHandler,
16
                     const property_list &propList = {});
17
18
      explicit queue(const device &syclDevice, const property_list &propList = {});
19
20
      explicit queue(const device &syclDevice, const async_handler &asyncHandler,
21
                     const property_list &propList = {});
22
23
      template <typename DeviceSelector>
24
      explicit queue(const context &syclContext,
25
                     const DeviceSelector &deviceSelector,
26
                     const property_list &propList = {});
27
28
      template <typename DeviceSelector>
29
      explicit queue(const context &syclContext,
30
                     const DeviceSelector &deviceSelector,
31
                     const async_handler &asyncHandler,
32
                     const property_list &propList = {});
33
34
      explicit queue(const context &syclContext, const device &syclDevice,
35
                     const property_list &propList = {});
36
37
      explicit queue(const context &syclContext, const device &syclDevice,
38
                     const async_handler &asyncHandler,
39
                     const property_list &propList = {});
40
```

96

```
/* -- common interface members -- */
41
42
43
      /* -- property interface members -- */
44
45
      backend get_backend() const;
46
47
      context get_context() const;
48
49
      device get_device() const;
50
51
      bool is_host() const;
52
53
      bool is_in_order() const;
54
55
      template <info::queue param>
56
      typename info::param_traits<info::queue, param>::return_type get_info() const;
57
58
      template <typename BackendEnum, BackendEnum param>
59
      typename info::param_traits<BackendEnum, param>::return_type
60
      get_backend_info() const;
61
62
      template <typename T>
63
      event submit(T cgf);
64
65
      template <typename T>
66
      event submit(T cgf, const queue &secondaryQueue);
67
68
      void wait();
69
70
      void wait_and_throw();
71
72
      void throw_asynchronous();
73
74
      /* -- convenience shortcuts -- */
75
76
      template <typename KernelName, typename KernelType>
77
      event single_task(const KernelType &KernelFunc);
78
79
      template <typename KernelName, typename KernelType>
80
      event single_task(event DepEvent, const KernelType &KernelFunc);
81
82
      template <typename KernelName, typename KernelType>
83
      event single_task(const std::vector<event> &DepEvents,
84
                        const KernelType &KernelFunc);
85
86
      template <typename KernelName, typename KernelType, int Dims>
87
      event parallel_for(range<Dims> NumWorkItems, const KernelType &KernelFunc);
88
89
      template <typename KernelName, typename KernelType, int Dims>
90
      event parallel_for(range<Dims> NumWorkItems, event DepEvent,
91
                         const KernelType &KernelFunc);
92
93
      template <typename KernelName, typename KernelType, int Dims>
94
      event parallel_for(range<Dims> NumWorkItems,
95
                         const std::vector<event> &DepEvents,
```

| 96 | <pre>const KernelType &KernelFunc);</pre> |
|-----|--|
| 97 | |
| 98 | template <typename dims="" int="" kernelname,="" kerneltype,="" typename=""></typename> |
| 99 | <pre>event parallel_for(range<dims> NumWorkItems, id<dims> WorkItemOffset,</dims></dims></pre> |
| 100 | <pre>const KernelType &KernelFunc);</pre> |
| 101 | |
| 102 | template <typename dims="" int="" kernelname,="" kerneltype,="" typename=""></typename> |
| 103 | <pre>event parallel_for(range<dims> NumWorkItems, id<dims> WorkItemOffset,</dims></dims></pre> |
| 104 | <pre>event DepEvent, const KernelType &KernelFunc);</pre> |
| 105 | |
| 106 | template <typename dims="" int="" kernelname,="" kerneltype,="" typename=""></typename> |
| 107 | <pre>event parallel_for(range<dims> NumWorkItems, id<dims> WorkItemOffset,</dims></dims></pre> |
| 108 | <pre>const std::vector<event> &DepEvents,</event></pre> |
| 109 | <pre>const KernelType &KernelFunc);</pre> |
| 110 | |
| 111 | template <typename dims="" int="" kernelname,="" kerneltype,="" typename=""></typename> |
| 112 | <pre>event parallel_for(nd_range<dims> ExecutionRange, const KernelType &KernelFunc);</dims></pre> |
| 113 | |
| 114 | template <typename dims="" int="" kernelname,="" kerneltype,="" typename=""></typename> |
| 115 | <pre>event parallel_for(nd_range<dims> ExecutionRange, event DepEvent,</dims></pre> |
| 116 | <pre>const KernelType &KernelFunc);</pre> |
| 117 | |
| 118 | template <typename dims="" int="" kernelname,="" kerneltype,="" typename=""></typename> |
| 119 | <pre>event parallel_for(nd_range<dims> ExecutionRange,</dims></pre> |
| 120 | <pre>const std::vector<event> &DepEvents,</event></pre> |
| 121 | <pre>const KernelType &KernelFunc);</pre> |
| 122 | }; |
| 123 | } // namespace sycl |

| Constructor | Description |
|---|---|
| <pre>explicit queue(const property_list &propList = {})</pre> | Constructs a SYCL queue instance |
| | using the device constructed from the |
| | default_selector_v. Zero or more |
| | properties can be provided to the con- |
| | structed SYCL queue via an instance of |
| | property_list. |
| <pre>explicit queue(const async_handler &asyncHandler,</pre> | Constructs a SYCL queue instance with an |
| <pre>const property_list &propList = {})</pre> | async_handler using the device constructed |
| | from the default_selector_v. Zero or |
| | more properties can be provided to the con- |
| | structed SYCL queue via an instance of |
| | property_list. |
| <pre>template <typename deviceselector=""></typename></pre> | Constructs a SYCL queue instance using the |
| <pre>explicit queue(const DeviceSelector &</pre> | device returned by the device selector pro- |
| deviceSelector, | vided. Zero or more properties can be pro- |
| <pre>const property_list &propList = {})</pre> | vided to the constructed SYCL queue via an |
| | instance of property_list. |
| | Continued on next page |

Table 4.21: Constructors of the queue class.

| template <typename deviceselector="">Constructs a SYCL queue instance with an async_handler using the device returned by the device selector provided. Zero or more properties can be provided to the con- structed SYCL queue via an instance of property_list.explicit queue(const device &syclDevice,Constructs a SYCL queue instance using</typename> |
|---|
| explicit queue(const DeviceSelector & deviceSelector,async_handler using the device returned by the device selector provided. Zero or more properties can be provided to the con- structed SYCL queue via an instance of property_list.explicit queue(const device &syclDevice,Constructs a SYCL queue instance using |
| deviceSelector,by the device selector provided. Zero orconst async_handler &asyncHandler,more properties can be provided to the con-const property_list &propList = {})structed SYCL queue via an instance ofexplicit queue(const device &syclDevice,Constructs a SYCL queue instance using |
| const async_handler &asyncHandler, const property_list &propList = {})more properties can be provided to the con- structed SYCL queue via an instance of property_list.explicit queue(const device &syclDevice,Constructs a SYCL queue instance using |
| <pre>const property_list &propList = {}) explicit queue(const device &syclDevice,</pre> |
| property_list. explicit queue(const device &syclDevice, Constructs a SYCL queue instance using |
| explicit queue(const device &syclDevice, Constructs a SYCL queue instance using |
| |
| <pre>const property_list &propList = {})</pre> the syclDevice provided. Zero or more |
| properties can be provided to the con- |
| structed SYCL queue via an instance of |
| property_list. |
| explicit queue(const device &syclDevice, Constructs a SYCL queue instance with an |
| const async_handler & asyncHandler, async_handler using the syclDevice pro- |
| const property_list &propList = {}) vided. Zero or more properties can be pro- |
| vided to the constructed SYCL queue via an |
| instance of property_list. |
| template <typename deviceselector=""> Constructs a SYCL queue instance that is</typename> |
| explicit quede(const context asycicontext, associated with the sycicontext provided, using the device returned by the device set |
| const preventies and a single device selector and a single device retained by the device se- |
| with the orrect invalid chiect orrect er |
| ror code if suclContext does not en- |
| cansulate the SYCL device returned by |
| deviceSelector. Zero or more properties |
| can be provided to the constructed SYCL |
| queue via an instance of property list. |
| template <typename deviceselector=""> Constructs a SYCL queue instance with</typename> |
| explicit queue(const context &syclContext, an async_handler that is associated with |
| const DeviceSelector & deviceSelector, the syclContext provided, using the de- |
| const async_handler &asyncHandler, vice returned by the device selector pro- |
| <pre>const property_list &propList = {}) vided. Must throw an exception with the</pre> |
| <pre>errc::invalid_object_error error code</pre> |
| if syclContext does not encapsulate the |
| SYCL device returned by deviceSelector |
| . Zero or more properties can be provided to |
| the constructed SYCL queue via an instance |
| of property_list. |
| explicit queue(const context &syclContext, Constructs a SYCL queue instance using |
| const device &syclDevice, the syclDevice provided, and associ- |
| <pre>const property_list &propList = {}) ated with the syclContext provided.</pre> |
| Must throw an exception with the errc |
| ::invalid_object_error error code if |
| syclContext does not encapsulate the |
| SYCL device syclDevice. Zero or more |
| properties can be provided to the con- |
| structed SYCL queue via an instance of |
| property_11st. |

Table 4.21: Constructors of the queue class.

| Constructor | Description |
|---|--|
| <pre>explicit queue(const context &syclContext,</pre> | Constructs a SYCL queue instance with an |
| const device &syclDevice, | async_handler using the syclDevice pro- |
| <pre>const async_handler &asyncHandler,</pre> | vided, and associated with the syclContext |
| <pre>const property_list &propList = {})</pre> | provided. Must throw an exception |
| | with the errc::invalid_object_error er- |
| | ror code if syclContext does not encapsu- |
| | late the SYCL device syclDevice. Zero |
| | or more properties can be provided to the |
| | constructed SYCL queue via an instance of |
| | property_list. |
| | End of table |

Table 4.21: Constructors of the queue class.

| Member function | Description |
|--|--|
| <pre>backend get_backend()const</pre> | Returns the a backend identifying the SYCL |
| | backend associated with this queue. |
| <pre>context get_context ()const</pre> | Returns the SYCL queue's context. Reports |
| | errors using SYCL exception classes. The |
| | value returned must be equal to that returned |
| | <pre>by get_info<info::queue::context>().</info::queue::context></pre> |
| <pre>device get_device ()const</pre> | Returns the SYCL device the queue is as- |
| | sociated with. Reports errors using SYCL |
| | exception classes. The value returned must |
| | be equal to that returned by get_info <info< td=""></info<> |
| | ::queue::devices>(). |
| <pre>bool is_host()const</pre> | Returns true if the backend associated with |
| | this SYCL queue is a SYCL host backend. |
| <pre>bool is_in_order()const</pre> | Returns true if the SYCL queue was cre- |
| | ated with the in_order property. Equiv- |
| | alent to has_property <property::queue::< td=""></property::queue::<> |
| | in_order>(). |
| <pre>void wait()</pre> | Performs a blocking wait for the comple- |
| | tion of all enqueued tasks in the queue. |
| | Synchronous errors will be reported through |
| | SYCL exceptions. |
| | Continued on next page |

Table 4.22: Member functions for queue class.

| Member function | Description |
|--|--|
| <pre>void wait_and_throw ()</pre> | Performs a blocking wait for the comple- |
| | tion of all enqueued tasks in the queue. |
| | Synchronous errors will be reported through |
| | SYCL exceptions. Any unconsumed asyn- |
| | chronous errors will be passed to the |
| | async_handler associated with the queue |
| | or enclosing context. If no user defined |
| | async_handler is associated with the queue |
| | or enclosing context, then an implemen- |
| | tation defined default async_handler is |
| | called to handle any errors, as described in |
| | 4.15.1.2. |
| <pre>void throw_asynchronous ()</pre> | Checks to see if any unconsumed asyn- |
| | chronous errors have been produced by the |
| | queue and if so reports them by passing them |
| | to the async_handler associated with the |
| | queue or enclosing context. If no user de- |
| | fined async_handler is associated with the |
| | queue or enclosing context, then an imple- |
| | mentation defined default async_handler |
| | is called to handle any errors, as described |
| | in 4.15.1.2. |
| <pre>template <info::queue param=""></info::queue></pre> | Queries this SYCL queue for information |
| <pre>typename info::param_traits</pre> | requested by the template parameter param |
| <info::queue, param="">::return_type</info::queue,> | . Specializations of info::param_traits |
| get_info ()const | must be defined in accordance with the info |
| | parameters in Table 4.23 to facilitate return- |
| | ing the type associated with the param pa- |
| | rameter. |
| template <typename t=""></typename> | Submit a command group function object to |
| <pre>event submit(T cgf)</pre> | the queue, in order to be scheduled for exe- |
| | cution on the device. |
| template <typename t=""></typename> | Submit a command group function object to |
| <pre>event submit(T cgf,</pre> | the queue, in order to be scheduled for ex- |
| queue & secondaryQueue) | ecution on the device. On a kernel error, |
| | this command group function object, is then |
| | scheduled for execution on the secondary |
| | queue. Returns an event, which corresponds |
| | to the queue the command group function |
| | object is being enqueued on. |
| | Continued on next page |

Table 4.22: Member functions for queue class.

| Member function | Description |
|---|--|
| <pre>template <typename backendenum="" backendenum,="" param=""></typename></pre> | Queries this SYCL queue for SYCL back- |
| <pre>typename info::param_traits<backendenum, param="">::</backendenum,></pre> | end-specific information requested by the |
| return_type | template parameter param. BackendEnum |
| <pre>get_backend_info()const</pre> | can be any enum class type specified |
| | by the SYCL backend specification of a |
| | supported SYCL backend named accord- |
| | ing to the convention info:: <backend_name< td=""></backend_name<> |
| | >::queue and param must be a valid |
| | enumeration of that enum class. Spe- |
| | cializations of info::param_traits must |
| | be defined for BackendEnum in accor- |
| | dance with the SYCL backend specifica- |
| | tion. Must throw an exception with the |
| | errc::invalid_object_error error code if |
| | the SYCL backend that corresponds with |
| | BackendEnum is different from the SYCL |
| | backend that is associated with this queue. |
| <pre>template <typename kernelname,="" kerneltype="" typename=""></typename></pre> | Defines and invokes a SYCL kernel function |
| <pre>event single_task(const KernelType &kernelFunc)</pre> | as a lambda function or a named function |
| | object type. The kernel function is submit- |
| | ted to the queue, in order to be scheduled for |
| | execution on the device. |
| template <typename kernelname,="" kerneltype="" typename=""></typename> | Defines and invokes a SYCL kernel function |
| event single_task(| as a lambda function or a named function |
| event DepEvent, const Kernellype &kernelFunc) | tad to the guesse in order to be scheduled |
| | for execution on the device once the event |
| | specified by DenEvent has completed |
| tomplate stymoname KernelName tymoname KernelTymes | Defines and invokes a SVCL kernel function |
| event single task(| as a lambda function or a named function |
| const_std::vector <event> &DenEvents</event> | object type The kernel function is submit- |
| const KernelType &kernelFunc) | ted to the queue in order to be scheduled |
| | for execution on the device once every event |
| | specified by DepEvents has completed. |
| template <typename kernelname,="" kerneltvpe.<="" td="" typename=""><td>Defines and invokes a SYCL kernel func-</td></typename> | Defines and invokes a SYCL kernel func- |
| int dimensions> | tion as a lambda function or a named func- |
| <pre>event parallel_for(</pre> | tion object type, for the specified range and |
| <pre>range<dimensions> numWorkItems,</dimensions></pre> | given an id or item for indexing in the in- |
| <pre>const KernelType &kernelFunc)</pre> | dexing space defined by range. The kernel |
| | function is submitted to the queue, in order |
| | to be scheduled for execution on the device. |
| | Continued on next page |

Table 4.22: Member functions for queue class.

| Member function | Description |
|---|--|
| <pre>template <typename kernelname,="" kerneltype,<="" pre="" typename=""></typename></pre> | Defines and invokes a SYCL kernel func- |
| int dimensions> | tion as a lambda function or a named func- |
| <pre>event parallel_for(</pre> | tion object type, for the specified range and |
| <pre>range<dimensions> numWorkItems, event DepEvent,</dimensions></pre> | given an id or item for indexing in the index- |
| <pre>const KernelType &kernelFunc)</pre> | ing space defined by range. The kernel func- |
| | tion is submitted to the queue, in order to be |
| | scheduled for execution on the device once |
| | the event specified by DepEvent has com- |
| | pleted. |
| <pre>template <typename kernelname,="" kerneltype,<="" pre="" typename=""></typename></pre> | Defines and invokes a SYCL kernel func- |
| int dimensions> | tion as a lambda function or a named func- |
| event parallel_for(| tion object type, for the specified range and |
| range <dimensions> numWorkItems,</dimensions> | given an id or item for indexing in the in- |
| const std::vector <event> &DepEvents,</event> | dexing space defined by range. The kernel |
| const Kernellype &kernelFunc) | to be scheduled for execution on the device |
| | once every event specified by DenEvents |
| | has completed |
| template <typename kernelname="" kerneltype<="" td="" typename=""><td>Defines and invokes a SYCL kernel func-</td></typename> | Defines and invokes a SYCL kernel func- |
| int dimensions> | tion as a lambda function or a named func- |
| event parallel for(| tion object type, for the specified range and |
| <pre>range<dimensions> numWorkItems,</dimensions></pre> | given an id or item for indexing in the in- |
| <pre>id<dimensions> workItemOffset,</dimensions></pre> | dexing space defined by range. The kernel |
| <pre>const KernelType &kernelFunc)</pre> | function is submitted to the queue, in order |
| | to be scheduled for execution. |
| <pre>template <typename kernelname,="" kerneltype,<="" pre="" typename=""></typename></pre> | Defines and invokes a SYCL kernel func- |
| int dimensions> | tion as a lambda function or a named func- |
| <pre>event parallel_for(</pre> | tion object type, for the specified range and |
| <pre>range<dimensions> numWorkItems,</dimensions></pre> | given an id or item for indexing in the index- |
| <pre>id<dimensions> workItemOffset,</dimensions></pre> | ing space defined by range. The kernel func- |
| event DepEvent, | tion is submitted to the queue, in order to be |
| const KernelType &kernelFunc) | scheduled for execution on the device once |
| | the event specified by DepEvent has com- |
| | Defines and involves a SVCL formal funa |
| int dimensions | tion as a lambda function or a named func- |
| avent parallel for (| tion object type, for the specified range and |
| range dimensions numWork Items | given an id or item for indexing in the in- |
| id <dimensions> workItemOffset.</dimensions> | dexing space defined by range. The kernel |
| <pre>const std::vector<event> &DepEvents.</event></pre> | function is submitted to the aueue. in order |
| const KernelType &kernelFunc) | to be scheduled for execution on the device |
| | once every event specified by DepEvents |
| | has completed. |
| | Continued on next page |

Table 4.22: Member functions for queue class.

| Member function | Description |
|---|--|
| <pre>template <typename kernelname,="" kerneltype,<="" pre="" typename=""></typename></pre> | Defines and invokes a SYCL kernel function |
| <pre>int dimensions></pre> | as a lambda function or a named function |
| <pre>event parallel_for(</pre> | object type, for the specified nd-range and |
| <pre>nd_range<dimensions> executionRange,</dimensions></pre> | given an nd-item for indexing in the index- |
| <pre>const KernelType &kernelFunc)</pre> | ing space defined by the nd-range. The ker- |
| | nel function is submitted to the queue, in or- |
| | der to be scheduled for execution. |
| <pre>template <typename kernelname,="" kerneltype,<="" pre="" typename=""></typename></pre> | Defines and invokes a SYCL kernel function |
| <pre>int dimensions></pre> | as a lambda function or a named function |
| <pre>event parallel_for(</pre> | object type, for the specified nd-range and |
| <pre>nd_range<dimensions> executionRange,</dimensions></pre> | given an nd-item for indexing in the index- |
| <pre>event DepEvent, const &KernelType kernelFunc)</pre> | ing space defined by the nd-range. The ker- |
| | nel function is submitted to the queue, in or- |
| | der to be scheduled for execution on the de- |
| | vice once the event specified by DepEvent |
| | has completed. |
| <pre>template <typename kernelname,="" kerneltype,<="" pre="" typename=""></typename></pre> | Defines and invokes a SYCL kernel func- |
| int dimensions> | tion as a lambda function or a named func- |
| <pre>event parallel_for(</pre> | tion object type, for the specified nd-range |
| <pre>nd_range<dimensions> executionRange,</dimensions></pre> | and given an nd-item for indexing in the in- |
| <pre>const std::vector<event> &DepEvents,</event></pre> | dexing space defined by the nd-range. The |
| <pre>const KernelType &kernelFunc)</pre> | kernel function is submitted to the queue, |
| | in order to be scheduled for execution on |
| | the device once every event specified by |
| | DepEvents has completed. |
| | End of table |

Table 4.22: Member functions for queue class.

4.6.5.2 Queue information descriptors

A queue can be queried for information using the get_info member function of the queue class, specifying one of the info parameters enumerated in info::queue. Every queue (including a host queue) must produce a valid value for each info parameter. The possible values for each info parameter and any restriction are defined in the specification of the SYCL backend associated with the queue. All info parameters in info::queue are specified in Table 4.23 and the synopsis for info::queue is described in appendix A.4.

| Queue Descriptors | Return type | Description |
|--------------------------------|-------------|--|
| info::queue::context | context | Returns the SYCL context associated with |
| | | this SYCL queue. |
| <pre>info::queue::device</pre> | device | Returns the SYCL device associated with this |
| | | SYCL queue. |
| | | End of table |
| | | |

Table 4.23: Queue information descriptors.

4.6.5.3 Queue properties

The properties that can be provided when constructing the SYCL queue class are describe in Table 4.24.

| Property | Description |
|--|---|
| <pre>property::queue::enable_profiling</pre> | The enable_profiling property adds the |
| | requirement that the SYCL runtime must |
| | capture profiling information for the com- |
| | mand groups that are submitted from |
| | this SYCL queue and provide said in- |
| | formation via the SYCL event class |
| | <pre>get_profiling_info member function, if</pre> |
| | the associated SYCL device has aspect:: |
| | queue_profiling. |
| <pre>property::queue::in_order</pre> | The in_order property adds the require- |
| | ment that the SYCL queue provides in-order |
| | semantics where tasks are executed in the or- |
| | der in which they are submitted. Tasks sub- |
| | mitted in this fashion can be viewed as hav- |
| | ing an implicit dependence on the previously |
| | submitted operation. |
| | End of table |

Table 4.24: Properties supported by the SYCL queue class.

The constructors of the queue property classes are listed in Table 4.25.

| Constructor | Description |
|--|---|
| <pre>property::queue::enable_profiling::enable_profiling</pre> | Constructs a SYCL enable_profiling |
| 0 | property instance. |
| <pre>property::queue::in_order::in_order()</pre> | Constructs a SYCL in_order property in- |
| | stance. |
| | End of table |

Table 4.25: Constructors of the queue property classes.

4.6.5.4 Queue error handling

Queue errors come in two forms:

- Synchronous Errors are those that we would expect to be reported directly at the point of waiting on an event, and hence waiting for a queue to complete, as well as any immediate errors reported by enqueuing work onto a queue. Such errors are reported through C++ exceptions.
- Asynchronous errors are those that are produced or detected after associated host API calls have returned (so can't be thrown as exceptions by the API call), and that are handled by an async_handler through which the errors are reported. Handling of asynchronous errors from a queue occurs at specific times, as described by 4.15.

Note that if there are asynchronous errors to be processed when a queue is destructed, the handler is called and this might delay or block the destruction, according to the behavior of the handler.

4.6.6 Event class

An event in SYCL is an object that represents the status of an operation that is being executed by the SYCL runtime.

Typically in SYCL, data dependency and execution order is handled implicitly by the SYCL runtime. However, in some circumstances developers want fine grain control of the execution, or want to retrieve properties of a command that is running.

A SYCL event maps to a single SYCL backend object when available. Note that, although an event represents the status of a particular operation, the dependencies of a certain event can be used to keep track of multiple steps required to synchronize said operation.

A SYCL event is returned by the submission of a command group. The dependencies of the event returned via the submission of the command group are the implementation-defined commands associated with the command group execution.

The SYCL event class provides the common reference semantics (see Section 4.5.3).

The constructors and member functions of the SYCL event class are listed in Tables 4.26 and 4.27, respectively. The additional common special member functions and common member functions are listed in Tables 4.1 and 4.2, respectively.

```
namespace sycl {
 1
2
3
    class event {
 4
     public:
5
      event();
6
7
      /* -- common interface members -- */
8
9
      backend get_backend() const;
10
11
      bool is_host() const;
12
13
      std::vector<event> get_wait_list();
14
15
      void wait();
16
17
      static void wait(const std::vector<event> &eventList);
18
19
      void wait_and_throw();
20
21
      static void wait_and_throw(const std::vector<event> &eventList);
22
23
      template <info::event param>
24
      typename info::param_traits<info::event, param>::return_type
25
      get_info() const;
26
      template <typename BackendEnum, BackendEnum param>
27
28
      typename info::param_traits<BackendEnum, param>::return_type
29
      get_backend_info() const;
30
31
      template <info::event_profiling param>
```

32 typename info::param_traits<info::event_profiling, param>::return_type

- 33 get_profiling_info() const;
- 34 };
- 35

36 } // namespace sycl

| event () Constructs a host event that is immediately ready. The event has no dependencies and no associated commands. Waiting on this event will return immediately. End of table | Constructor | Description |
|---|-------------|---|
| ready. The event has no dependencies and no associated commands. Waiting on this event will return immediately. End of table | event () | Constructs a host event that is immediately |
| no associated commands. Waiting on this event will return immediately. End of table | | ready. The event has no dependencies and |
| event will return immediately. End of table | | no associated commands. Waiting on this |
| End of table | | event will return immediately. |
| | | End of table |

Table 4.26: Constructors of the event class.

| Member function | Description |
|---|---|
| <pre>backend get_backend()const</pre> | Returns the a backend identifying the SYCL |
| | backend associated with this event. |
| <pre>bool is_host()const</pre> | Returns true if this SYCL event is a host |
| | event. |
| <pre>std::vector<event> get_wait_list()</event></pre> | Return the list of events that this event waits |
| | for in the dependence graph. Only direct de- |
| | pendencies are returned, and not transitive |
| | dependencies that direct dependencies wait |
| | on. Whether already completed events are |
| | included in the returned list is implementa- |
| | tion defined. |
| void wait() | Wait for the event and the command associ- |
| | ated with it to complete. |
| <pre>void wait_and_throw()</pre> | Wait for the event and the command associ- |
| | ated with it to complete. |
| | Any unconsumed asynchronous errors from |
| | any context that the event was waiting |
| | on executions from will be passed to the |
| | async_handler associated with the con- |
| | text. If no user defined async_handler is |
| | associated with the context, then an imple- |
| | mentation defined default async_handler |
| | is called to handle any errors, as described |
| | in 4.15.1.2. |
| static void wait(| Synchronously wait on a list of events. |
| <pre>const std::vector<event> &eventList)</event></pre> | |
| | Continued on next page |

Table 4.27: Member functions for the event class.

| Member function | Description |
|---|--|
| <pre>static void wait_and_throw(</pre> | Synchronously wait on a list of events. |
| <pre>const std::vector<event> &eventList)</event></pre> | Any unconsumed asynchronous errors from |
| | any context that the event was waiting |
| | on executions from will be passed to the |
| | async_handler associated with the con- |
| | text. If no user defined async_handler is |
| | associated with the context, then an imple- |
| | mentation defined default async_handler |
| | is called to handle any errors, as described |
| | in 4.15.1.2. |
| <pre>template <info::event param=""></info::event></pre> | Queries this SYCL event for information |
| <pre>typename info::param_traits</pre> | requested by the template parameter param |
| <info::event, param="">::return_type</info::event,> | . Specializations of info::param_traits |
| <pre>get_info()const</pre> | must be defined in accordance with the info |
| | parameters in Table 4.28 to facilitate return- |
| | ing the type associated with the param pa- |
| | rameter. |
| <pre>template <typename backendenum="" backendenum,="" param=""></typename></pre> | Queries this SYCL event for SYCL back- |
| <pre>typename info::param_traits<backendenum, param="">::</backendenum,></pre> | end-specific information requested by the |
| return_type | template parameter param. BackendEnum |
| <pre>get_backend_info()const</pre> | can be any enum class type specified |
| | by the SYCL backend specification of a |
| | supported SYCL backend named accord- |
| | ing to the convention info:: <backend_name< td=""></backend_name<> |
| | >::event and param must be a valid |
| | enumeration of that enum class. Spe- |
| | cializations of info::param_traits must |
| | be defined for BackendEnum in accor- |
| | dance with the SYCL backend specifica- |
| | tion. Must throw an exception with the |
| | errc::invalid_object_error error code if |
| | the SYCL backend that corresponds with |
| | BackendEnum is different from the SYCL |
| | backend that is associated with this event. |
| | Continued on next page |

Table 4.27: Member functions for the event class.

| Member function | Description |
|--|---|
| <pre>Member function template <info::event_profiling param=""> typename info::param_traits <info::event_profiling, param="">::return_type get_profiling_info ()const</info::event_profiling,></info::event_profiling></pre> | Description Queries this SYCL event for profiling in- formation requested by the parameter param . If the requested profiling information is unavailable when get_profiling_info is called due to incompletion of command groups associated with the event, then the call to get_profiling_info will block until the requested profiling information is avail- able. An example is asking for info:: event_profiling::command_end when the associated command group has yet to fin- ish execution. Calls to get_profiling_info must throw an exception with the errc ::invalid_object_error error code if the SYCL queue which submitted the command group this SYCL event is associated with was not constructed with the property:: queue::enable_profiling property. Spe- cializations of info::param_traits must be defined in accordance with the info parame- |
| | ters in Table 4.29 to facilitate returning the |
| | type associated with the param parameter. |
| | End of table |

Table 4.27: Member functions for the event class.

4.6.6.1 Event information and profiling descriptors

An event can be queried for information using the get_info member function of the event class, specifying one of the info parameters enumerated in info::event. Every event (including a host event) must produce a valid value for each info parameter. The possible values for each info parameter and any restrictions are defined in the specification of the SYCL backend associated with the event. All info parameters in info::event are specified in Table 4.28 and the synopsis for info::event is described in appendix A.6.

| Event Descriptors | Return type | Description |
|--------------------------|---------------------|---|
| info::event:: | info:: | Returns the event status of the command group |
| command_execution_status | event_command_statu | s associated with this SYCL event. |
| | | |
| | | End of table |

Table 4.28: Event class information descriptors.

An event can be queried for profiling information using the get_profiling_info member function of the event class, specifying one of the profiling info parameters enumerated in info::event_profiling. Every event (including a host event) must produce a valid value for each info parameter. The possible values for each info parameter and any restrictions are defined in the specification of the SYCL backend associated with the event. All info parameters in info::event_profiling are specified in Table 4.29 and the synopsis for info::event_profiling is described in appendix A.6.
| Event information profiling descrip- | Return type | Description |
|--------------------------------------|-------------|---|
| tor | | |
| <pre>info::event_profiling::</pre> | uint64_t | Returns an implementation defined 64-bit |
| command_submit | | value describing the time in nanoseconds when |
| | | the associated command group was submitted. |
| <pre>info::event_profiling::</pre> | uint64_t | Returns an implementation defined 64-bit |
| command_start | | value describing the time in nanoseconds when |
| | | the associated command group started execut- |
| | | ing. |
| <pre>info::event_profiling::</pre> | uint64_t | Returns an implementation defined 64-bit |
| command_end | | value describing the time in nanoseconds when |
| | | the associated command group finished execut- |
| | | ing. |
| | | End of table |

Table 4.29: Profiling information descriptors for the SYCL event class.

4.7 Data access and storage in SYCL

In SYCL, data storage and access are handled by separate classes. Buffers and images handle storage and ownership of the data, whereas accessors handle access to the data. Buffers and images in SYCL can be bound to more than one device or context, including across different SYCL backends. They also handle ownership of the data, while allowing exception handling for blocking and non-blocking data transfers. Accessors manage data transfers between the host and all of the devices in the system, as well as tracking of data dependencies.

4.7.1 Host allocation

A SYCL runtime may need to allocate temporary objects on the host to handle some operations (such as copying data from one context to another). Allocation on the host is managed using an allocator object, following the standard C++ allocator class definition. The default allocator for memory objects is implementation defined, but the user can supply their own allocator class.

```
1 {
2 buffer<int, 1, UserDefinedAllocator<int> > b(d);
3 }
```

When an allocator returns a nullptr, the runtime cannot allocate data on the host. Note that in this case the runtime will raise an error if it requires host memory but it is not available (e.g when moving data across SYCL backend contexts).

The definition of allocators extends the current functionality of SYCL, ensuring that users can define allocator functions for specific hardware or certain complex shared memory mechanisms (e.g. NUMA), and improves interoperability with STL-based libraries (e.g. Intel's TBB provides an allocator).

4.7.1.1 Default allocators

A default allocator is always defined by the implementation, and it is guaranteed to return non-nullptr and new memory positions every call. The default allocator for const buffers will remove the const-ness of the type (therefore, the default allocator for a buffer of type "const int" will be an Allocator<int>). This implies that

host accessors will not synchronize with the pointer given by the user in the buffer/image constructor, but will use the memory returned by the Allocator itself for that purpose. The user can implement an allocator that returns the same address as the one passed in the buffer constructor, but it is the responsibility of the user to handle the potential race conditions.

| Allocators | Description |
|------------------|---|
| buffer_allocator | It is the default buffer allocator used by the |
| | runtime, when no allocator is defined by the |
| | user. |
| image_allocator | It is the default allocator used by the run- |
| | time for the SYCL unsampled_image and |
| | <pre>sampled_image classes when no allocator is</pre> |
| | provided by the user. The image_allocator |
| | is required allocate in elements of std:: |
| | byte. |
| | End of table |

Table 4.30: SYCL Default Allocators.

See Section 4.7.5 for details on manual host-device synchronization.

4.7.2 Buffers

The **buffer** class defines a shared array of one, two or three dimensions that can be used by the SYCL kernel and has to be accessed using accessor classes. Buffers are templated on both the type of their data, and the number of dimensions that the data is stored and accessed through.

A **buffer** does not map to only one underlying backend object, and all **SYCL** backend memory objects may be temporary for use within a command group on a specific device. Note that if no source data is provided for a buffer, the buffer uses uninitialized memory for performance reasons. So it is up to the programmer to explicitly construct the objects in this case if required.

More generally, since the value type of a buffer is required to be trivially copyable, there is no constructor or destructor called in any case.

A SYCL buffer can construct an instance of a SYCL buffer that reinterprets the original SYCL buffer with a different type, dimensionality and range using the member function reinterpret. The reinterpreted SYCL buffer that is constructed must behave as though it were a copy of the SYCL buffer that constructed it (see see 4.5.3) with the exception that the type, dimensionality and range of the reinterpreted SYCL buffer must reflect the type, dimensionality and range specified when calling the reinterpret member function. By extension of this, the class member types value_type, reference and const_reference, and the member functions get_range and get_count of the reinterpreted SYCL buffer must reflect the new type, dimensionality and range. The data that the original SYCL buffer and the reinterpreted SYCL buffer may alter to reflect the new type, dimensionality and range. It is important to note that a reinterpreted SYCL buffer is a copy of the original SYCL buffer. Constructing more than one SYCL buffer managing the same host pointer is still undefined behavior.

The SYCL buffer class template provides the common reference semantics (see Section 4.5.3).

4.7.2.1 Buffer interface

The constructors and member functions of the SYCL **buffer** class template are listed in Tables 4.31 and 4.32, respectively. The additional common special member functions and common member functions are listed in Tables 4.1 and 4.2, respectively.

Each constructor takes as the last parameter an optional SYCL property_list to provide properties to the SYCL buffer.

The SYCL buffer class template takes a template parameter AllocatorT for specifying an allocator which is used by the SYCL runtime when allocating temporary memory on the host. If no template argument is provided, then the default allocator for the SYCL buffer class buffer_allocator will be used (see 4.7.1.1).

```
1 namespace sycl {
2 namespace property {
3 namespace buffer {
4 class use_host_ptr {
5
     public:
6
        use_host_ptr() = default;
7 };
8
9
   class use_mutex {
10
     public:
11
        use_mutex(std::mutex &mutexRef);
12
13
        std::mutex *get_mutex_ptr() const;
14
   };
15
   class context_bound {
16
17
     public:
18
        context_bound(context boundContext);
19
20
        context get_context() const;
21 };
22 } // namespace buffer
23 } // namespace property
24
25
   template <typename T, int dimensions = 1,</pre>
26
              typename AllocatorT = sycl::buffer_allocator>
27
    class buffer {
28
     public:
     using value_type = T;
29
30
     using reference = value_type &;
31
     using const_reference = const value_type &;
32
     using allocator_type = AllocatorT;
33
34
     buffer(const range<dimensions> &bufferRange,
35
             const property_list &propList = {});
36
37
     buffer(const range<dimensions> &bufferRange, AllocatorT allocator,
38
             const property_list &propList = {});
39
40
      buffer(T *hostData, const range<dimensions> &bufferRange,
41
             const property_list &propList = {});
42
```

```
43
      buffer(T *hostData, const range<dimensions> &bufferRange,
44
             AllocatorT allocator, const property_list &propList = {});
45
46
      buffer(const T *hostData, const range<dimensions> &bufferRange,
47
             const property_list &propList = {});
48
49
      buffer(const T *hostData, const range<dimensions> &bufferRange,
50
             AllocatorT allocator, const property_list &propList = {});
51
52
      buffer(const std::shared_ptr<T> &hostData,
53
             const range<dimensions> &bufferRange, AllocatorT allocator,
54
             const property_list &propList = {});
55
56
      buffer(const std::shared_ptr<T> &hostData,
57
             const range<dimensions> &bufferRange,
58
             const property_list &propList = {});
59
60
      template <class InputIterator>
61
      buffer<T, 1>(InputIterator first, InputIterator last, AllocatorT allocator,
62
                   const property_list &propList = {});
63
      template <class InputIterator>
64
65
      buffer<T, 1>(InputIterator first, InputIterator last,
66
                   const property_list &propList = {});
67
      buffer(buffer<T, dimensions, AllocatorT> b, const id<dimensions> &baseIndex,
68
69
             const range<dimensions> &subRange);
70
71
      /* -- common interface members -- */
72
73
      /* -- property interface members -- */
74
75
      range<dimensions> get_range() const;
76
77
      size_t get_count() const;
78
79
      size_t get_size() const;
80
81
      AllocatorT get_allocator() const;
82
83
      template <access::mode mode, access::target target = access::target::global_buffer>
84
      accessor<T, dimensions, mode, target> get_access(
85
          handler &commandGroupHandler);
86
87
      template <access::mode mode>
88
      accessor<T, dimensions, mode, access::target::host_buffer> get_access();
89
90
      template <access::mode mode, access::target target = access::target::global_buffer>
91
      accessor<T, dimensions, mode, target> get_access(
92
          handler &commandGroupHandler, range<dimensions> accessRange,
93
          id<dimensions> accessOffset = {});
94
95
      template <access::mode mode>
96
      accessor<T, dimensions, mode, access::target::host_buffer> get_access(
97
        range<dimensions> accessRange, id<dimensions> accessOffset = {});
```

```
98
99
       template<typename... Ts>
100
       auto get_access(Ts...);
101
102
       template<typename... Ts>
103
       auto get_host_access(Ts...);
104
105
       template <typename Destination = std::nullptr_t>
106
       void set_final_data(Destination finalData = nullptr);
107
108
       void set_write_back(bool flag = true);
109
110
       bool is_sub_buffer() const;
111
112
       template <typename ReinterpretT, int ReinterpretDim>
113
       buffer<ReinterpretT, ReinterpretDim, AllocatorT>
114
      reinterpret(range<ReinterpretDim> reinterpretRange) const;
115
116
      // Only available when ReinterpretDim == 1
117
      // or when (ReinterpretDim == dimensions) &&
118
                 (sizeof(ReinterpretT) == sizeof(T))
      11
119
       template <typename ReinterpretT, int ReinterpretDim = dimensions>
120
       buffer<ReinterpretT, ReinterpretDim, AllocatorT>
121
      reinterpret() const;
122 };
123
124 // Deduction guides
125 template <class InputIterator, class AllocatorT>
126
    buffer(InputIterator, InputIterator, AllocatorT, const property_list & = {})
127
         -> buffer<typename std::iterator_traits<InputIterator>::value_type, 1,
128
                  AllocatorT>;
129 template <class InputIterator>
130 buffer(InputIterator, InputIterator, const property_list & = {})
        -> buffer<typename std::iterator_traits<InputIterator>::value_type, 1>;
131
132 template <class T, int dimensions, class AllocatorT>
133 buffer(const T *, const range<dimensions> &, AllocatorT,
134
           const property_list & = {})
135
         -> buffer<T, dimensions, AllocatorT>;
136 template <class T, int dimensions>
137
    buffer(const T *, const range<dimensions> &, const property_list & = {})
         -> buffer<T, dimensions>;
138
139
140 } // namespace sycl
```

| Constructor | Description |
|---|--|
| <pre>buffer(const range<dimensions> & bufferRange, const property_list &propList = {})</dimensions></pre> | Construct a SYCL buffer instance with uninitialized memory. The constructed SYCL buffer will use a default constructed AllocatorT when allocating memory on the host. The range of the constructed SYCL buffer is specified by the bufferRange pa- rameter provided. Data is not written back to the host on destruction of the buffer unless the buffer has a valid non-null pointer specified via the member function set_final_data(). Zero or more properties can be provided to the constructed SYCL buffer via an instance of property_list. |
| <pre>buffer(const range<dimensions> & bufferRange, AllocatorT allocator, const property_list &propList = {})</dimensions></pre> | Construct a SYCL buffer instance with uninitialized memory. The constructed SYCL buffer will use the allocator pa- rameter provided when allocating mem- ory on the host. The range of the con- structed SYCL buffer is specified by the bufferRange parameter provided. Data is not written back to the host on destruc- tion of the buffer unless the buffer has a valid non-null pointer specified via the member function set_final_data(). Zero or more properties can be provided to the constructed SYCL buffer via an instance of property_list. |
| <pre>buffer(T* hostData, const range<dimensions> & bufferRange, const property_list &propList = {})</dimensions></pre> | Construct a SYCL buffer instance with the hostData parameter provided. The buffer is initialized with the memory specified by hostData. The ownership of this memory is given to the constructed SYCL buffer for the duration of its lifetime. The constructed SYCL buffer will use a default constructed AllocatorT when allocating memory on the host. The range of the constructed SYCL buffer is specified by the bufferRange pa- rameter provided. Zero or more properties can be provided to the constructed SYCL buffer via an instance of property_list. |

Table 4.31: Constructors of the **buffer** class.

| Constructor | Description |
|--|--|
| <pre>buffer(T* hostData,</pre> | Construct a SYCL buffer instance with the |
| <pre>const range<dimensions> & bufferRange,</dimensions></pre> | hostData parameter provided. The buffer |
| AllocatorT allocator, | is initialized with the memory specified by |
| <pre>const property_list &propList = {})</pre> | hostData. The ownership of this memory is |
| | given to the constructed SYCL buffer for |
| | the duration of its lifetime. The constructed |
| | SYCL buffer will use the allocator pa- |
| | rameter provided when allocating mem- |
| | ory on the host. The range of the con- |
| | structed SYCL buffer is specified by the |
| | bufferRange parameter provided. Zero or |
| | more properties can be provided to the con- |
| | structed SYCL buffer via an instance of |
| | property_list. |
| <pre>buffer(const T* hostData,</pre> | Construct a SYCL buffer instance with the |
| <pre>const range<dimensions> & bufferRange,</dimensions></pre> | hostData parameter provided. The owner- |
| <pre>const property_list &propList = {})</pre> | ship of this memory is given to the con- |
| | structed SYCL buffer for the duration of its |
| | lifetime. |
| | The constructed SYCL buffer will use a de- |
| | fault constructed AllocatorT when allocat- |
| | ing memory on the host. |
| | The host address is const T, so the host |
| | accesses can be read-only. However, the |
| | typename T is not const so the device ac- |
| | cesses can be both read and write accesses. |
| | Since the hostData is const, this buffer is |
| | only initialized with this memory and there |
| | is no write back after its destruction, unless |
| | the buffer has another valid non-null final |
| | data address specified via the member func- |
| | tion set_final_data() after construction of |
| | The range of the constructed SVCL buffer |
| | is specified by the by Company permitter |
| | ns specified by the bufferkange parameter |
| | Zaro or more properties can be provided |
| | to the constructed SVCL buffer via on in |
| | stance of property list |
| | Stance of property_fist. |
| | Continued on next page |

| Constructor | Description |
|--|--|
| <pre>buffer(const T* hostData,</pre> | Construct a SYCL buffer instance with the |
| <pre>const range<dimensions> & bufferRange,</dimensions></pre> | hostData parameter provided. The owner- |
| AllocatorT allocator, | ship of this memory is given to the con- |
| <pre>const property_list &propList = {})</pre> | structed SYCL buffer for the duration of its |
| | lifetime. |
| | The constructed SYCL buffer will use the |
| | allocator parameter provided when allo- |
| | cating memory on the host. |
| | The host address is const T, so the host |
| | accesses can be read-only. However, the |
| | typename T is not const so the device ac- |
| | cesses can be both read and write accesses. |
| | Since, the hostData is const, this buffer is |
| | only initialized with this memory and there |
| | is no write back after its destruction, unless |
| | the buffer has another valid non-null final |
| | data address specified via the member func- |
| | tion set_final_data() after construction of |
| | the buffer. |
| | The range of the constructed SYCL buffer |
| | is specified by the bufferRange parameter |
| | provided. |
| | Zero or more properties can be provided |
| | to the constructed SYCL buffer via an in- |
| | stance of property_list. |
| <pre>buffer(const std::shared_ptr<t> &hostData,</t></pre> | Construct a SYCL buffer instance with the |
| <pre>const range<dimensions> & bufferRange,</dimensions></pre> | hostData parameter provided. The owner- |
| <pre>const property_list &propList = {})</pre> | ship of this memory is given to the con- |
| | structed SYCL buffer for the duration of its |
| | lifetime. The constructed SYCL buffer will |
| | use a default constructed AllocatorT when |
| | allocating memory on the host. The range of |
| | the constructed SYCL buffer is specified by |
| | the bufferRange parameter provided. Zero |
| | or more properties can be provided to the |
| | constructed SYCL buffer via an instance of |
| | property_list. |
| | Continued on next page |

| Constructor | Description |
|---|---|
| <pre>buffer(const std::shared_ptr<void> &hostData,</void></pre> | Construct a SYCL buffer instance with the |
| <pre>const range<dimensions> & bufferRange,</dimensions></pre> | hostData parameter provided. The owner- |
| AllocatorT allocator, | ship of this memory is given to the con- |
| <pre>const property_list &propList = {})</pre> | structed SYCL buffer for the duration of its |
| | lifetime. The constructed SYCL buffer will |
| | use the allocator parameter provided when |
| | allocating memory on the host. The range of |
| | the constructed SYCL buffer is specified by |
| | the bufferRange parameter provided. Zero |
| | or more properties can be provided to the |
| | constructed SYCL buffer via an instance of |
| | property_list. |
| <pre>template <typename inputiterator=""></typename></pre> | Create a new allocated 1D buffer initialized |
| <pre>buffer(InputIterator first, InputIterator last,</pre> | from the given elements ranging from first |
| <pre>const property_list &propList = {})</pre> | up to one before last. The data is copied |
| | to an intermediate memory position by the |
| | runtime. Data is not written back to the same |
| | iterator set provided. However, if the buffer |
| | has a valid non-const iterator specified |
| | via the member function set_final_data |
| | (), data will be copied back to that iter- |
| | ator. The constructed SYCL buffer will |
| | use a default constructed AllocatorT when |
| | allocating memory on the host. Zero or |
| | more properties can be provided to the con- |
| | structed SYCL buffer via an instance of |
| | property_list. |
| template <typename inputiterator=""></typename> | Create a new allocated ID buffer initial- |
| buffer(InputIterator first, InputIterator last, | ized from the given elements ranging from |
| AllocatorT allocator = {}, | first up to one before last. The data |
| <pre>const property_list &propList = {})</pre> | is copied to an intermediate memory po- |
| | sition by the runtime. Data is not writ- |
| | ten back to the same iterator set provided. |
| | However, if the buffer has a valid non- |
| | const iterator specified via the member func- |
| | tion set_final_data(), data will be copied |
| | back to that herator. The constructed SYCL |
| | provided when allocating memory on the |
| | host. Zero or more properties can be are |
| | vided to the constructed SVCL buffer vie |
| | an instance of property list |
| | Continued on payt page |
| | Continued on next page |

| Constructor | Description |
|--|--|
| <pre>buffer(buffer<t, allocatort="" dimensions,=""> &b,</t,></pre> | Create a new sub-buffer without allocation |
| <pre>const id<dimensions> & baseIndex,</dimensions></pre> | to have separate accessors later. b is the |
| <pre>const range<dimensions> & subRange)</dimensions></pre> | buffer with the real data, which must not |
| | be a sub-buffer. baseIndex specifies the |
| | origin of the sub-buffer inside the buffer b |
| | . subRange specifies the size of the sub- |
| | buffer. The sum of baseIndex and subRange |
| | in any dimension must not exceed the par- |
| | ent buffer (b) size (bufferRange) in that di- |
| | mension, and an exception with the errc:: |
| | <pre>invalid_object_error error code must be</pre> |
| | thrown if violated. |
| | The offset and range specified by baseIndex |
| | and subRange together must represent a con- |
| | tiguous region of the original SYCL buffer. |
| | If a non-contiguous region of a buffer |
| | is requested when constructing a sub- |
| | buffer, then an exception with the errc:: |
| | invalid_object_error error code must be |
| | thrown. |
| | The origin (based on baseIndex) of the sub- |
| | buffer being constructed must be a multi- |
| | ple of the memory base address alignment |
| | of each SYCL device that is executed on, |
| | otherwise the SYCL funtime must throw an |
| | asynchronous exception with the errc:: |
| | This value is retrievable via the SVCI |
| | device class info query information |
| | uevice class into query info::device |
| | throw exception with the error |
| | invalid chiect error error code if h |
| | is already a sub-buffer |
| | Find of table |

| Member function | Description |
|--|--|
| <pre>range<dimensions> get_range()const</dimensions></pre> | Return a range object representing the size |
| | of the buffer in terms of number of elements |
| | in each dimension as passed to the construc- |
| | tor. |
| <pre>size_t get_count()const</pre> | Returns the total number of elements in the |
| | buffer. Equal to get_range()[0] * * |
| | <pre>get_range()[dimensions-1].</pre> |
| <pre>size_t get_size()const</pre> | Returns the size of the buffer storage in |
| | bytes. Equal to get_count()*sizeof(T). |
| | Continued on next page |

Table 4.32: Member functions for the **buffer** class.

| Member function | Description |
|---|--|
| AllocatorT get_allocator()const | Returns the allocator provided to the buffer. |
| <pre>template<access::mode access::target="" mode,="" target="</pre"></access::mode></pre> | Returns a valid accessor to the buffer with |
| <pre>access::target::global_buffer></pre> | the specified access mode and target in the |
| <pre>accessor<t, dimensions,="" mode,="" target=""></t,></pre> | command group buffer. The value of target |
| <pre>get_access(handler &commandGroupHandler)</pre> | can be access::target::global_buffer or |
| | <pre>access::constant_buffer.</pre> |
| <pre>template<access::mode mode=""></access::mode></pre> | Returns a valid host accessor to the buffer |
| <pre>accessor<t, access::target::<="" dimensions,="" mode,="" pre=""></t,></pre> | with the specified access mode and target. |
| host_buffer> | |
| <pre>get_access()</pre> | |
| <pre>template<access::mode access::target="" mode,="" target="</pre"></access::mode></pre> | Returns a valid accessor to the buffer |
| <pre>access::target::global_buffer></pre> | with the specified access mode and tar- |
| <pre>accessor<t, dimensions,="" mode,="" target=""></t,></pre> | get in the command group buffer. Only |
| <pre>get_access(handler &commandGroupHandler, range<</pre> | the values starting from the given offset |
| <pre>dimensions> accessRange, id<dimensions> accessOffset</dimensions></pre> | and up to the given range are guaran- |
| = {}) | teed to be updated. The value of target |
| | <pre>can be access::target::global_buffer or</pre> |
| | <pre>access::constant_buffer.</pre> |
| <pre>template<access::mode mode=""></access::mode></pre> | Returns a valid host accessor to the buffer |
| <pre>accessor<t, access::target::<="" dimensions,="" mode,="" pre=""></t,></pre> | with the specified access mode and target. |
| host_buffer> | Only the values starting from the given off- |
| <pre>get_access(range<dimensions> accessRange, id<</dimensions></pre> | set and up to the given range are guaranteed |
| <pre>dimensions> accessOffset = {})</pre> | to be updated. The value of target can only |
| | <pre>be access::target::host_buffer.</pre> |
| <pre>template<typename ts=""></typename></pre> | Returns a valid accessor as if constructed |
| <pre>auto get_access(Ts args)</pre> | via passing the buffer and all provided |
| | arguments to the SYCL accessor. |
| | |
| | Possible implementation: |
| | <pre>return accessor { *this, args };</pre> |
| <pre>template<typename ts=""></typename></pre> | Returns a valid host_accessor as if |
| <pre>auto get_host_access(Ts args)</pre> | constructed via passing the buffer and |
| | all provided arguments to the SYCL |
| | host_accessor. |
| | |
| | Possible implementation: |
| | <pre>return host_accessor { *this, args</pre> |
| | }; |
| | Continued on next page |

Table 4.32: Member functions for the **buffer** class.

| Member function | Description |
|---|--|
| <pre>template <typename destination="std::nullptr_t"></typename></pre> | The finalData points to where the outcome |
| <pre>void set_final_data(Destination finalData =</pre> | of all the buffer processing is going to be |
| nullptr) | copied to at destruction time, if the buffer |
| | was involved with a write accessor. |
| | Destination can be either an output iterator |
| | or a std::weak_ptr <t>.</t> |
| | Note that a raw pointer is a special case |
| | of output iterator and thus defines the host |
| | memory to which the result is to be copied. |
| | In the case of a weak pointer, the output is |
| | not updated if the weak pointer has expired. |
| | If Destination 18 std::nullptr_t, then the |
| | This means an for stion allows demonstration |
| Vold Set_write_back(bool flag = true) | forcing or canceling the write back of the |
| | data of a buffer on destruction according to |
| | the value of flag |
| | Forcing the write-back is similar to what |
| | happens during a normal write-back as de- |
| | scribed in Section 4.7.2.3 and 4.7.4. |
| | If there is nowhere to write-back, using this |
| | function does not have any effect. |
| <pre>bool is_sub_buffer()const</pre> | Returns true if this SYCL buffer is a sub- |
| | buffer, otherwise returns false. |
| <pre>template <typename int="" reinterpretdim="" reinterprett,=""></typename></pre> | Creates and returns a reinterpreted |
| <pre>buffer<reinterprett, allocatort="" reinterpretdim,=""></reinterprett,></pre> | SYCL buffer with the type specified |
| <pre>reinterpret(range<reinterpretdim></reinterpretdim></pre> | by ReinterpretT, dimensions specified by |
| reinterpretRange)const | ReinterpretDim and range specified by |
| | reinterpretRange. The buffer object being |
| | reinterpreted can be a SYCL sub-buffer |
| | that was created from a SYCL buffer and |
| | must throw exception with the errc:: |
| | invalid_object_error error code if the |
| | total size in bytes represented by the type |
| | (or sub buffer) does not equal the total |
| | size in bytes represented by the type and |
| | range of this SYCL buffer (or sub-buffer) |
| | Reinterpreting a sub-buffer provides a |
| | reinterpreted view of the sub-buffer only |
| | and does not change the offset or size of |
| | the sub-buffer view (in bytes) relative to the |
| | parent buffer. |
| | Continued on next page |

Table 4.32: Member functions for the **buffer** class.

| Member function | Description |
|---|---|
| <pre>template <typename int="" pre="" reinterpretdim<="" reinterprett,=""></typename></pre> | Creates and returns a reinterpreted |
| = dimensions> | SYCL buffer with the type specified |
| <pre>buffer<reinterprett, allocatort="" reinterpretdim,=""></reinterprett,></pre> | by ReinterpretT and dimensions spec- |
| reinterpret()const | ified by ReinterpretDim. Only valid |
| | when (ReinterpretDim == 1) or when |
| | ((ReinterpretDim == dimensions)&& (|
| | <pre>sizeof(ReinterpretT)== sizeof(T))).</pre> |
| | The buffer object being reinterpreted can be |
| | a SYCL sub-buffer that was created from a |
| | SYCL buffer and must throw an exception |
| | with the errc::invalid_object_error er- |
| | ror code if the total size in bytes represented |
| | by the type and range of the reinterpreted |
| | SYCL buffer (or sub-buffer) does not equal |
| | the total size in bytes represented by the |
| | type and range of this SYCL buffer (or |
| | sub-buffer). Reinterpreting a sub-buffer |
| | provides a reinterpreted view of the sub- |
| | buffer only, and does not change the offset |
| | or size of the sub-buffer view (in bytes) |
| | relative to the parent buffer . |
| | End of table |

Table 4.32: Member functions for the **buffer** class.

4.7.2.2 Buffer properties

The properties that can be provided when constructing the SYCL buffer class are describe in Table 4.33.

| Property | Description |
|---|---|
| <pre>property::buffer::use_host_ptr</pre> | The use_host_ptr property adds the re- |
| | quirement that the SYCL runtime must not |
| | allocate any memory for the SYCL buffer |
| | and instead uses the provided host pointer |
| | directly. This prevents the SYCL runtime |
| | from allocating additional temporary stor- |
| | age on the host. |
| <pre>property::buffer::use_mutex</pre> | The use_mutex property is valid for |
| | the SYCL buffer, unsampled_image and |
| | <pre>sampled_image classes. The property adds</pre> |
| | the requirement that the memory which is |
| | owned by the SYCL buffer can be shared |
| | with the application via a std::mutex pro- |
| | vided to the property. The mutex m is locked |
| | by the runtime whenever the data is in use |
| | and unlocked otherwise. Data is synchro- |
| | nized with hostData, when the mutex is un- |
| | locked by the runtime. |
| | Continued on next page |

Table 4.33: Properties supported by the SYCL buffer class.CHAPTER 4. SYCL PROGRAMMING INTERFACE

| Property | Description |
|--|--|
| <pre>property::buffer::context_bound</pre> | The context_bound property adds the re- |
| | quirement that the SYCL buffer can only |
| | be associated with a single SYCL context |
| | that is provided to the property. |
| | End of table |

Table 4.33: Properties supported by the SYCL **buffer** class.

The constructors and special member functions of the buffer property classes are listed in Tables 4.34 and 4.35 respectively.

| Constructor | Description |
|--|--|
| <pre>property::buffer::use_host_ptr::use_host_ptr()</pre> | Constructs a SYCL use_host_ptr property |
| | instance. |
| <pre>property::buffer::use_mutex::use_mutex(std::mutex &</pre> | Constructs a SYCL use_mutex property in- |
| mutexRef) | stance with a reference to mutexRef param- |
| | eter provided. |
| <pre>property::buffer::context_bound::context_bound(</pre> | Constructs a SYCL context_bound prop- |
| <pre>context boundContext)</pre> | erty instance with a copy of a SYCL |
| | context. |
| | End of table |

Table 4.34: Constructors of the **buffer property** classes.

| Member function | Description |
|---|---------------------------------------|
| <pre>std::mutex *property::buffer::use_mutex::</pre> | Returns the std::mutex which was |
| <pre>get_mutex_ptr()const</pre> | specified when constructing this SYCL |
| | use_mutex property. |
| <pre>context property::buffer::context_bound::get_context</pre> | Returns the context which was spec- |
| ()const | ified when constructing this SYCL |
| | context_bound property. |
| | End of table |

Table 4.35: Member functions of the **buffer property** classes.

4.7.2.3 Buffer synchronization rules

Buffers are reference-counted. When a buffer value is constructed from another buffer, the two values reference the same buffer and a reference count is incremented. When a buffer value is destroyed, the reference count is decremented. Only when there are no more buffer values that reference a specific buffer is the actual buffer destroyed and the buffer destruction behavior defined below is followed.

If any error occurs on buffer destruction, it is reported via the associated queue's asynchronous error handling mechanism.

The basic rule for the blocking behavior of a buffer destructor is that it blocks if there is some data to write back because a write accessor on it has been created, or if the buffer was constructed with attached host memory and is still in use.

More precisely:

- 1. A buffer can be constructed with just a size and using the default buffer allocator. The memory management for this type of buffer is entirely handled by the SYCL system. The destructor for this type of buffer does not need to block, even if work on the buffer has not completed. Instead, the SYCL system frees any storage required for the buffer asynchronously when it is no longer in use in queues. The initial contents of the buffer are unspecified.
- 2. A buffer can be constructed with associated host memory and a default buffer allocator. The buffer will use this host memory for its full lifetime, but the contents of this host memory are unspecified for the lifetime of the buffer. If the host memory is modified by the host, or mapped to another buffer or image during the lifetime of this buffer, then the results are undefined. The initial contents of the buffer will be the contents of the host memory at the time of construction.

When the buffer is destroyed, the destructor will block until all work in queues on the buffer have completed, then copy the contents of the buffer back to the host memory (if required) and then return.

(a) If the type of the host data is const, then the buffer is read-only; only read accessors are allowed on the buffer and no-copy-back to host memory is performed (although the host memory must still be kept available for use by SYCL). When using the default buffer allocator, the const-ness of the type will be removed in order to allow host allocation of memory, which will allow temporary host copies of the data by the SYCL runtime, for example for speeding up host accesses.

When the buffer is destroyed, the destructor will block until all work in queues on the buffer have completed and then return, as there is no copy of data back to host.

(b) If the type of the host data is not const but the pointer to host data is const, then the read-only restriction applies only on host and not on device accesses.

When the buffer is destroyed, the destructor will block until all work in queues on the buffer have completed.

3. A buffer can be constructed using a shared_ptr to host data. This pointer is shared between the SYCL application and the runtime. In order to allow synchronization between the application and the runtime a mutex is used which will be locked by the runtime whenever the data is in use, and unlocked when it is no longer needed.

The shared_ptr reference counting is used in order to prevent destroying the buffer host data prematurely. If the shared_ptr is deleted from the user application before buffer destruction, the buffer can continue securely because the pointer hasn't been destroyed yet. It will not copy data back to the host before destruction, however, as the application side has already deleted its copy.

Note that since there is an implicit conversion of a std::unique_ptr to a std::shared_ptr, a std:: unique_ptr can also be used to pass the ownership to the SYCL runtime.

4. A buffer can be constructed from a pair of iterator values. In this case, the buffer construction will copy the data from the data range defined by the iterator pair. The destructor will not copy back any data and does not need to block.

If set_final_data() is used to change where to write the data back to, then the destructor of the buffer will block if a write accessor on it has been created.

A sub-buffer object can be created which is a sub-range reference to a base buffer. This sub-buffer can be used

to create accessors to the base buffer, which have access to the range specified at time of construction of the sub-buffer. Sub-buffers cannot be created from sub-buffers, but only from a base buffer which is not already a sub-buffer.

Sub-buffers must be constructed from a contiguous region of memory in a buffer. This requirement is potentially non-intuitive when working with buffers that have dimensionality larger than one, but maps to one-dimensional SYCL backend native allocations without performance cost due to index mapping computation. For example:

```
1 buffer<int,2> parent_buffer { range<2>{ 8,8 } }; // Create 2-d buffer with 8x8 ints
2
3 // OK: Contiguous region from middle of buffer
4 buffer<int,2> sub_buf1 { parent_buffer, /*offset*/ range<2>{ 2,0 }, /*size*/ range<2>{ 2,8 } };
5
6 // invalid_object_error exception: Non-contiguous regions of 2-d buffer
7 buffer<int,2> sub_buf2 { parent_buffer, /*offset*/ range<2>{ 2,0 }, /*size*/ range<2>{ 2,2 } };
8 buffer<int,2> sub_buf3 { parent_buffer, /*offset*/ range<2>{ 2,2 }, /*size*/ range<2>{ 2,6 } };
9
10 // invalid_object_error exception: Out-of-bounds size
11 buffer<int,2> sub_buf4 { parent_buffer, /*offset*/ range<2>{ 2,2 }, /*size*/ range<2>{ 2,8 } };
```

4.7.3 Images

The classes unsampled_image (Table 4.36) and sampled_image (Table 4.38) define shared image data of one, two or three dimensions, that can be used by kernels in queues and have to be accessed using accessor classes with image accessor modes.

The constructors and member functions of the SYCL unsampled_image and sampled_image class templates are listed in Tables 4.36, 4.37, 4.38 and 4.39, respectively. The additional common special member functions and common member functions are listed in Tables 4.1 and 4.2, respectively.

Where relevant, it is the responsibility of the user to ensure that the format of the data matches the format described by image_format.

The allocator template parameter of the SYCL unsampled_image and sampled_image classes can be any allocator type including a custom allocator, however it must allocate in units of std::byte.

For any image that is constructed with the range (r1, r2, r3) with an element type size in bytes of *s*, the image row pitch and image slice pitch should be calculated as follows:

$$r1 \cdot s$$
 (4.1)

$$r1 \cdot r2 \cdot s \tag{4.2}$$

The SYCL unsampled_image and sampled_image class templates provide the common reference semantics (see Section 4.5.3).

4.7.3.1 Unsampled image interface

Each constructor of the unsampled_image takes an image_format to describe the data layout of the image data.

Each constructor additionally takes as the last parameter an optional SYCL property_list to provide properties to the SYCL unsampled_image.

The SYCL unsampled_image class template takes a template parameter AllocatorT for specifying an allocator which is used by the SYCL runtime when allocating temporary memory on the host. If no template argument is provided, the default allocator for the SYCL unsampled_image class image_allocator is used (see Section 4.7.1.1).

```
1 namespace sycl {
    enum class image_format : unsigned int {
2
     r8g8b8a8_unorm,
3
4
     r16g16b16a16_unorm,
5
      r8g8b8a8_sint,
     r16g16b16a16_sint,
6
7
     r32b32g32a32_sint,
8
     r8g8b8a8_uint,
9
     r16g16b16a16_uint,
10
     r32b32g32a32_uint,
11
     r16b16g16a16_sfloat,
12
     r32g32b32a32_sfloat,
13
     b8g8r8a8_unorm,
14 };
15
    template <int dimensions = 1, typename AllocatorT = sycl::image_allocator>
16
17
    class unsampled_image {
18
    public:
19
     unsampled_image(image_format format, const range<dimensions> &rangeRef,
20
                      const property_list &propList = {});
21
22
     unsampled_image(image_format format, const range<dimensions> &rangeRef,
23
                      AllocatorT allocator, const property_list &propList = {});
24
25
      /* Available only when: dimensions > 1 */
26
      unsampled_image(image_format format, const range<dimensions> &rangeRef,
27
                      const range<dimensions -1> &pitch,
28
                      const property_list &propList = {});
29
30
      /* Available only when: dimensions > 1 */
31
      unsampled_image(image_format format, const range<dimensions> &rangeRef,
32
                      const range<dimensions -1> &pitch, AllocatorT allocator,
33
                      const property_list &propList = {});
34
35
     unsampled_image(void *hostPointer, image_format format,
36
                      const range<dimensions> &rangeRef,
37
                      const property_list &propList = {});
38
39
      unsampled_image(void *hostPointer, image_format format,
40
                      const range<dimensions> &rangeRef, AllocatorT allocator,
41
                      const property_list &propList = {});
42
      /* Available only when: dimensions > 1 */
43
44
      unsampled_image(void *hostPointer, image_format format,
45
                      const range<dimensions> &rangeRef,
46
                      const range<dimensions -1> &pitch,
```

```
47
                       const property_list &propList = {});
48
49
       /* Available only when: dimensions > 1 */
50
       unsampled_image(void *hostPointer, image_format format,
51
                       const range<dimensions> &rangeRef,
52
                       const range<dimensions -1> &pitch, AllocatorT allocator,
53
                       const property_list &propList = {});
54
55
       unsampled_image(std::shared_ptr<void> &hostPointer, image_format format,
56
                       const range<dimensions> &rangeRef,
57
                       const property_list &propList = {});
58
59
       unsampled_image(std::shared_ptr<void> &hostPointer, image_format format,
60
                       const range<dimensions> &rangeRef, AllocatorT allocator,
61
                       const property_list &propList = {});
62
63
       /* Available only when: dimensions > 1 */
64
       unsampled_image(std::shared_ptr<void> &hostPointer, image_format format,
65
                       const range<dimensions> &rangeRef,
66
                       const range<dimensions -1> &pitch,
67
                       const property_list &propList = {});
68
69
       /* Available only when: dimensions > 1 */
70
       unsampled_image(std::shared_ptr<void> &hostPointer, image_format format,
71
                       const range<dimensions> &rangeRef,
72
                       const range<dimensions -1> &pitch, AllocatorT allocator,
73
                       const property_list &propList = {});
74
75
       /* -- common interface members -- */
76
       /* -- property interface members -- */
77
78
79
       range<dimensions> get_range() const;
80
81
       /* Available only when: dimensions > 1 */
82
       range<dimensions - 1> get_pitch() const;
83
84
       size_t get_count() const;
85
86
       size_t get_size() const;
87
88
       AllocatorT get_allocator() const;
89
90
       template <typename dataT, access::mode accessMode>
91
       accessor<dataT, dimensions, accessMode, access::target::unsampled_image>
92
       get_access(handler & commandGroupHandler);
93
94
       template <typename dataT, access::mode accessMode>
95
       accessor<dataT, dimensions, accessMode, access::target::host_unsampled_image>
96
       get_access();
97
98
       template <typename Destination = std::nullptr_t>
99
       void set_final_data(Destination finalData = std::nullptr);
100
101
       void set_write_back(bool flag = true);
```

102 };

103 } // namespace sycl

| Constructor | Description |
|--|--|
| <pre>unsampled_image(image_format format,</pre> | Construct a SYCL unsampled_image in- |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | stance with uninitialized memory. The con- |
| <pre>const property_list &propList = {})</pre> | structed SYCL unsampled_image will use a |
| | default constructed AllocatorT when allo- |
| | cating memory on the host. The element size |
| | of the constructed SYCL unsampled_image |
| | will be derived from the format pa- |
| | rameter. The range of the constructed |
| | SYCL unsampled_image is specified by the |
| | rangeRef parameter provided. The pitch |
| | of the constructed SYCL unsampled_image |
| | will be the default size determined by the |
| | SYCL runtime. Unless the member func- |
| | tion set_final_data() is called with a valid |
| | non-null pointer, there will be no write |
| | back on destruction. Zero or more prop- |
| | erties can be provided to the constructed |
| | SYCL unsampled_image via an instance of |
| | property_list. |
| <pre>unsampled_image(image_format format,</pre> | Construct a SYCL unsampled_image in- |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | stance with uninitialized memory. The con- |
| AllocatorT allocator, | structed SYCL unsampled_image will use |
| <pre>const property_list &propList = {})</pre> | the allocator parameter provided when |
| | allocating memory on the host. The |
| | element size of the constructed SYCL |
| | unsampled_image will be derived from |
| | the format parameter. The range of |
| | the constructed SYCL unsampled_image is |
| | specified by the rangeker parameter pro- |
| | vided. The pitch of the constructed STCL |
| | datarmined by the SVCL runtime. Unless |
| | the member function set final data() is |
| | called with a valid non null pointer there |
| | will be no write back on destruction. Zero |
| | or more properties can be provided to the |
| | constructed SYCL unsampled image via an |
| | instance of property list |
| | Continued on next page |

Table 4.36: Constructors of the unsampled_image class template.

| Constructor | Description |
|---|--|
| <pre>unsampled_image(image_format format,</pre> | Available only when: dimensions > 1. |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | Construct a SYCL unsampled_image in- |
| <pre>const range<dimensions-1> &pitch,</dimensions-1></pre> | stance with uninitialized memory. The con- |
| <pre>const property_list &propList = {})</pre> | structed SYCL unsampled_image will use a |
| | default constructed AllocatorT when allo- |
| | cating memory on the host. The element size |
| | of the constructed SYCL unsampled_image |
| | will be derived from the format pa- |
| | rameter. The range of the constructed |
| | SYCL unsampled_image is specified by the |
| | rangeRef parameter provided. The pitch |
| | of the constructed SYCL unsampled_image |
| | will be the pitch parameter provided. Un- |
| | less the member function set_final_data |
| | () is called with a valid non-null pointer, |
| | there will be no write back on destruction. |
| | Zero or more properties can be provided to |
| | the constructed SYCL unsampled_image via |
| | an instance of property_list. |
| <pre>unsampled_image(image_format format,</pre> | Available only when: dimensions > 1 . |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | Construct a SYCL unsampled_image in- |
| <pre>const range<dimensions-1> &pitch,</dimensions-1></pre> | stance with uninitialized memory. The con- |
| AllocatorT allocator, | structed SYCL unsampled_image will use |
| <pre>const property_list &propList = {})</pre> | the allocator parameter provided when |
| | allocating memory on the host. The |
| | element size of the constructed SYCL |
| | the format nonumeter. The range of |
| | the constructed SVCI successful dimension |
| | the constructed STCL unsampled_image is |
| | vided. The pitch of the constructed SVCI |
| | vided. The pitch of the constructed STCL |
| | eter provided Unless the member func |
| | tion set final data() is called with a valid |
| | non-null pointer there will be no write |
| | how has pointer, there will be no wille back on destruction. Zero or more prop- |
| | erties can be provided to the constructed |
| | SYCL unsampled image via an instance of |
| | property list. |
| | Continued on next page |

| <pre>unsampled_image(void *hostPointer, image_format format, const property_list &propList = {})</pre> Construct a SYCL unsampled_image in- stance with the hostPointer parameter provided. The ownership of this mem- ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use a default con- structed AllocatorT when allocating memory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer, const range | Constructor | Description |
|--|--|---|
| <pre>image_format format, const property_list &propList = {}) stance with the hostPointer parameter provided. The ownership of this mem- ory is given to the constructed SYCL unsampled_image will use a default con- structed AllocatorT when allocating memory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image in- stance with the hostPointer parameter provided. The ownership of this mem- ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be derived from the format parame- the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be detrived form the format parame- the rangeRef pa</pre> | <pre>unsampled_image(void *hostPointer,</pre> | Construct a SYCL unsampled_image in- |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | <pre>image_format format,</pre> | stance with the hostPointer parameter |
| <pre>const property_list &propList = {}) ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use a default con- structed AllocatorT when allocating memory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image (void *hostPointer, image_format format, const range<dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {}) </dimensions></pre> | <pre>const range<dimensions> &rangeRef,</dimensions></pre> | provided. The ownership of this mem- |
| <pre>unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use a default con- structed AllocatorT when allocating memory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list. const range<dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {})</dimensions></pre> | <pre>const property_list &propList = {})</pre> | ory is given to the constructed SYCL |
| <pre>its lifetime. The constructed SYCL unsampled_image will use a default con- structed AllocatorT when allocating memory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> | | unsampled_image for the duration of |
| <pre>unsampled_image will use a default con- structed AllocatorT when allocating memory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more pro- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> | | its lifetime. The constructed SYCL |
| <pre>structed AllocatorT when allocating memory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> | | unsampled_image will use a default con- |
| <pre>memory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> | | structed AllocatorT when allocating |
| <pre>the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> | | memory on the host. The element size of |
| <pre>will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> | | the constructed SYCL unsampled_image |
| <pre>ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list. Construct a SYCL unsampled_image in- stance with the hostPointer parameter provided. The ownership of this mem- ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parameter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image</pre> | | will be derived from the format parame- |
| <pre>unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> unsampled_image(void *hostPointer, image_format format, const range <dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {}) Constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be default size determined by</dimensions> | | ter. The range of the constructed SYCL |
| <pre>rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list. unsampled_image(void *hostPointer, image_format format,</pre> | | unsampled_image is specified by the |
| <pre>of the constructed SYCL unsampled_image will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> Construct a SYCL unsampled_image in- stance with the hostPointer parameter provided. The ownership of this mem- ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | | rangeRef parameter provided. The pitch |
| <pre>will be the default size determined by the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> Construct a SYCL unsampled_image in- stance with the hostPointer parameter provided. The ownership of this mem- ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by the start size determined by the start size determined by the start size of the pitch of the constructed SYCL unsampled_image | | of the constructed SYCL unsampled_image |
| <pre>the SYCL runtime. Unless the member function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> Construct a SYCL unsampled_image in- stance with the hostPointer parameter provided. The ownership of this mem- ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | | will be the default size determined by |
| <pre>function set_final_data() is called with a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> unsampled_image(void *hostPointer, image_format format, const range <dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {}) const property_list &propList = {}) unsampled_image ior the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by</dimensions> | | the SYCL runtime. Unless the member |
| <pre>a valid non-null pointer, any memory allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> unsampled_image(void *hostPointer, image_format format, const range <dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {}) const property_list &propList = {}) unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by</dimensions> | | function set_final_data() is called with |
| <pre>allocated by the SYCL runtime is written back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.</pre> unsampled_image(void *hostPointer, image_format format, const range <dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {}) Construct a SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by</dimensions> | | a valid non-null pointer, any memory |
| back to hostPointer. Zero or more prop- erties can be provided to the constructed SYCL unsampled_image via an instance of property_list. unsampled_image(void *hostPointer, image_format format, const range <dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {}) unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by</dimensions> | | allocated by the SYCL runtime is written |
| erties can be provided to the constructed SYCL unsampled_image via an instance of property_list.unsampled_image(void *hostPointer, image_format format, const range <dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {})Construct a SYCL unsampled_image in- stance with the hostPointer parameter provided. The ownership of this mem- ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by</dimensions> | | back to hostPointer. Zero or more prop- |
| SYCL unsampled_image via an instance of property_list.unsampled_image(void *hostPointer, image_format format, const range <dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {})Construct a SYCL unsampled_image in- stance with the hostPointer parameter provided. The ownership of this mem- ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by</dimensions> | | erties can be provided to the constructed |
| property_list.unsampled_image(void *hostPointer, image_format format, const range <dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {})Construct a SYCL unsampled_image in- stance with the hostPointer parameter provided. The ownership of this mem- ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by</dimensions> | | SYCL unsampled_image via an instance of |
| <pre>unsampled_image(void *hostPointer, image_format format, const range<dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {})</dimensions></pre> Construct a SYCL unsampled_image in- stance with the hostPointer parameter provided. The ownership of this mem- ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | | property_list. |
| <pre>image_format format, const range<dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {}) some and the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by</dimensions></pre> | <pre>unsampled_image(void *hostPointer,</pre> | Construct a SYCL unsampled_image in- |
| <pre>const range<dimensions> &rangeRef, AllocatorT allocator, const property_list &propList = {})</dimensions></pre> provided. The ownership of this mem- ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | <pre>image_format format,</pre> | stance with the hostPointer parameter |
| AllocatorT allocator, const property_list &propList = {}) ory is given to the constructed SYCL unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | <pre>const range<dimensions> &rangeRef,</dimensions></pre> | provided. The ownership of this mem- |
| <pre>const property_list &propList = {}) unsampled_image for the duration of its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by</pre> | AllocatorT allocator, | ory is given to the constructed SYCL |
| its lifetime. The constructed SYCL unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | <pre>const property_list &propList = {})</pre> | unsampled_image for the duration of |
| unsampled_image will use the allocator parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | | its lifetime. The constructed SYCL |
| parameter provided when allocating mem- ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | | unsampled_image will use the allocator |
| ory on the host. The element size of the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | | parameter provided when allocating mem- |
| the constructed SYCL unsampled_image will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | | ory on the host. The element size of |
| will be derived from the format parame- ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | | the constructed SYCL unsampled_image |
| ter. The range of the constructed SYCL unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | | will be derived from the format parame- |
| unsampled_image is specified by the rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | | ter. The range of the constructed SYCL |
| rangeRef parameter provided. The pitch of the constructed SYCL unsampled_image will be the default size determined by | | unsampled_image is specified by the |
| of the constructed SYCL unsampled_image will be the default size determined by | | rangeRef parameter provided. The pitch |
| will be the default size determined by | | of the constructed SYCL unsampled_image |
| | | will be the default size determined by |
| the SYCL runtime. Unless the member | | the SYCL runtime. Unless the member |
| iunculon set_final_data() is called with | | runcuon set_rinal_data() is called with |
| a valid non-null pointer, any memory | | a value non-null pointer, any memory |
| allocated by the SYCL runtime is written | | have to heat Deinten. Zare on more such |
| Dack to nostPointer. Zero or more prop- | | ortion can be provided to the constructed |
| erties can be provided to the constructed | | SVCI uncompled image via an instance of |
| SICL unsampled_image via an instance of | | property list |
| property_list. | | Continued on next page |

| Constructor | Description |
|---|---|
| <pre>unsampled_image(void *hostPointer,</pre> | Available only when: dimensions > 1 |
| <pre>image_format format,</pre> | Construct a SYCL unsampled_image in- |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | stance with the hostPointer parameter |
| <pre>const range<dimensions-1> &pitch,</dimensions-1></pre> | provided. The ownership of this mem- |
| <pre>const property_list &propList = {})</pre> | ory is given to the constructed SYCL |
| | unsampled_image for the duration of |
| | its lifetime. The constructed SYCL |
| | unsampled_image will use a default con- |
| | structed AllocatorT when allocating |
| | memory on the host. The element size of |
| | the constructed SYCL unsampled_image |
| | will be derived from the format parame- |
| | ter. The range of the constructed SYCL |
| | unsampled_image is specified by the |
| | rangeRef parameter provided. The pitch |
| | of the constructed SYCL unsampled_image |
| | will be the pitch parameter provided. Un- |
| | less the member function set_final_data |
| | () is called with a valid non-null pointer, |
| | any memory allocated by the SYCL runtime |
| | is written back to hostPointer. Zero or |
| | more properties can be provided to the |
| | constructed SYCL unsampled_image via an |
| | instance of property_list. |
| unsampled_image(void *hostPointer, | Available only when: dimensions > 1 . |
| <pre>image_format format,</pre> | Construct a SYCL unsampled_image in- |
| const range <dimensions> &rangeRef,</dimensions> | stance with the nostPointer parameter |
| Const range <dimensions-i> &pitch,</dimensions-i> | provided. The ownership of this mem- |
| Allocatori allocator, | uncompled image for the duration of |
| const property_fist apropList = {}) | its lifetime The constructed SVCI |
| | uncompled image will use the allocator |
| | parameter provided when allocating mem |
| | ory on the host. The element size of |
| | the constructed SYCL unsampled image |
| | will be derived from the format parame- |
| | ter The range of the constructed SYCL |
| | unsampled image is specified by the |
| | rangeRef parameter provided. The pitch |
| | of the constructed SYCL unsampled image |
| | will be the pitch parameter provided. Un- |
| | less the member function set_final_data |
| | () is called with a valid non-null pointer, |
| | any memory allocated by the SYCL runtime |
| | is written back to hostPointer. Zero or |
| | more properties can be provided to the |
| | constructed SYCL unsampled_image via an |
| | instance of property_list. |
| | Continued on next page |

| Constructor | Description |
|---|---|
| <pre>unsampled_image(std::shared_ptr<void>& hostPointer,</void></pre> | Construct a SYCL unsampled_image in- |
| <pre>image_format format,</pre> | stance with the hostPointer parameter |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | provided. The ownership of this mem- |
| <pre>const property_list &propList = {})</pre> | ory is given to the constructed SYCL |
| | unsampled_image for the duration of |
| | its lifetime. The constructed SYCL |
| | unsampled_image will use a default con- |
| | structed AllocatorT when allocating |
| | memory on the host. The element size of |
| | the constructed SYCL unsampled_image |
| | will be derived from the format parame- |
| | ter. The range of the constructed SYCL |
| | unsampled_image is specified by the |
| | rangeRef parameter provided. The pitch |
| | of the constructed SYCL unsampled_image |
| | will be the default size determined by |
| | the SYCL runtime. Unless the member |
| | function set_final_data() is called with |
| | a valid non-null pointer, any memory |
| | allocated by the SYCL runtime is written |
| | back to hostPointer. Zero or more prop- |
| | erties can be provided to the constructed |
| | SYCL unsampled_image via an instance of |
| | property_list. |
| <pre>unsampled_image(std::shared_ptr<void>& hostPointer,</void></pre> | Construct a SYCL unsampled_image in- |
| image_format format, | stance with the hostPointer parameter |
| const range <dimensions> &rangeRef,</dimensions> | provided. The ownership of this mem- |
| Allocatori allocator, | ory is given to the constructed SYCL |
| const property_list &propList = {}) | unsampled_image for the duration of |
| | its metime. The constructed STCL |
| | parameter provided when allocating mem |
| | ory on the bost. The element size of |
| | the constructed SVCI unsampled image |
| | will be derived from the format parame- |
| | ter The range of the constructed SYCI |
| | unsampled image is specified by the |
| | rangeRef parameter provided The pitch |
| | of the constructed SYCL unsampled image |
| | will be the default size determined by |
| | the SYCL runtime. Unless the member |
| | function set_final_data() is called with |
| | a valid non-null pointer, any memory |
| | allocated by the SYCL runtime is written |
| | back to hostPointer. Zero or more prop- |
| | erties can be provided to the constructed |
| | SYCL unsampled_image via an instance of |
| | property_list. |
| | Continued on next page |

| Constructor | Description |
|---|--|
| <pre>unsampled_image(std::shared_ptr<void>& hostPointer,</void></pre> | Construct a SYCL unsampled_image in- |
| <pre>image_format format,</pre> | stance with the hostPointer parameter |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | provided. The ownership of this mem- |
| <pre>const range<dimensions-1> & pitch,</dimensions-1></pre> | ory is given to the constructed SYCL |
| <pre>const property_list &propList = {})</pre> | unsampled_image for the duration of |
| | its lifetime. The constructed SYCL |
| | unsampled_image will use a default con- |
| | structed AllocatorT when allocating |
| | memory on the host. The element size of |
| | the constructed SYCL unsampled_image |
| | will be derived from the format parame- |
| | ter. The range of the constructed SYCL |
| | unsampled_image is specified by the |
| | rangeRet parameter provided. The pitch |
| | of the constructed SYCL unsampled_image |
| | will be the pitch parameter provided. Un- |
| | () is called with a valid non null pointer |
| | any memory allocated by the SVCL runtime |
| | is written back to hostPointer. Zero or |
| | more properties can be provided to the |
| | constructed SYCL unsampled image via an |
| | instance of property list. |
| <pre>unsampled_image(std::shared_ptr<void>& hostPointer,</void></pre> | Construct a SYCL unsampled_image in- |
| image_format format, | stance with the hostPointer parameter |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | provided. The ownership of this mem- |
| <pre>const range<dimensions-1> & pitch,</dimensions-1></pre> | ory is given to the constructed SYCL |
| AllocatorT allocator, | unsampled_image for the duration of |
| <pre>const property_list &propList = {})</pre> | its lifetime. The constructed SYCL |
| | unsampled_image will use the allocator |
| | parameter provided when allocating mem- |
| | ory on the host. The element size of |
| | the constructed SYCL unsampled_image |
| | will be derived from the format parame- |
| | ter. The range of the constructed SYCL |
| | unsampled_image is specified by the |
| | of the constructed SVCL uncompled image |
| | will be the nitch parameter provided Un- |
| | less the member function set final data |
| | () is called with a valid non-null pointer |
| | any memory allocated by the SYCL runtime |
| | is written back to hostPointer. Zero or |
| | more properties can be provided to the |
| | constructed SYCL unsampled_image via an |
| | instance of property_list. |
| | End of table |

Table 4.36: Constructors of the unsampled_image class template.

| Member function | Description |
|--|---|
| <pre>range<dimensions> get_range()const</dimensions></pre> | Return a range object representing the size |
| | of the image in terms of the number of el- |
| | ements in each dimension as passed to the |
| | constructor. |
| <pre>range<dimensions-1> get_pitch()const</dimensions-1></pre> | Available only when: dimensions > 1. |
| | Return a range object representing the pitch |
| | of the image in bytes. |
| <pre>size_t get_count()const</pre> | Returns the total number of elements in the |
| | <pre>image. Equal to get_range()[0] * *</pre> |
| | <pre>get_range()[dimensions-1].</pre> |
| <pre>size_t get_size()const</pre> | Returns the size of the image storage |
| | in bytes. The number of bytes may |
| | be greater than get_count()*element size |
| | due to padding of elements, rows and slices |
| | of the image for efficient access. |
| AllocatorT get_allocator()const | Returns the allocator provided to the image. |
| <pre>template<typename access::mode="" accessmode="" datat,=""></typename></pre> | Returns a valid accessor to the image with |
| <pre>accessor<datat, access::<="" accessmode,="" dimensions,="" pre=""></datat,></pre> | the specified access mode and target. The |
| <pre>target::unsampled_image></pre> | only valid types for dataT are int4, uint4, |
| <pre>get_access(handler & commandGroupHandler)</pre> | float4 and half4. |
| <pre>template<typename access::mode="" accessmode="" datat,=""></typename></pre> | Returns a valid accessor to the image with |
| <pre>accessor<datat, access::<="" accessmode,="" dimensions,="" pre=""></datat,></pre> | the specified access mode and target. The |
| <pre>target::host_unsampled_image></pre> | only valid types for dataT are int4, uint4, |
| get_access() | float4 and half4. |
| <pre>template <typename destination="std::nullptr_t"></typename></pre> | The finalData point to where the output |
| <pre>void set_final_data(Destination finalData =</pre> | of all the image processing is going to be |
| nullptr) | copied to at destruction time, if the image |
| | was involved with a write accessor. |
| | Destination can be either an output iterator, |
| | <pre>a std::weak_ptr<t>.</t></pre> |
| | Note that a raw pointer is a special case |
| | of output iterator and thus defines the host |
| | memory to which the result is to be copied. |
| | In the case of a weak pointer, the output is |
| | not copied if the weak pointer has expired. |
| | If Destination is std::nullptr_t, then the |
| | copy back will not happen. |
| | |
| <pre>void set_write_back(bool flag = true)</pre> | This member function allows dynamically |
| | forcing or canceling the write-back of the |
| | to the value of Glass |
| | to the value of flag. |
| | Forcing the write-back is similar to what |
| | nappens during a normal write-back as de- |
| | scribed in Section 4.7.3.4 and 4.7.4. |
| | function does not have any effect |
| | Tunction does not have any effect. |

Table 4.37: Member functions of the unsampled_image class template.

4.7.3.2 Sampled image interface

Each constructor of the sampled_image class requires a pointer to the host data the image will sample, an image_format to describe the data layout and an image_sampler (Section 4.7.8) to describe how to sample the image data.

Each constructor additionally takes as the last parameter an optional SYCL property_list to provide properties to the SYCL sampled_image.

```
1 namespace sycl {
    enum class image_format : unsigned int {
 2
 3
      r8g8b8a8_unorm,
 4
      r16g16b16a16_unorm,
 5
      r8g8b8a8_sint,
 6
      r16g16b16a16_sint,
 7
      r32b32g32a32_sint,
 8
      r8g8b8a8_uint,
 9
      r16g16b16a16_uint,
10
      r32b32g32a32_uint,
11
      r16b16g16a16_sfloat,
12
      r32g32b32a32_sfloat,
13
      b8g8r8a8_unorm,
14 };
15
    template <int dimensions = 1, typename AllocatorT = sycl::image_allocator>
16
17
    class sampled_image {
18
     public:
19
      sampled_image(const void *hostPointer, image_format format,
20
                      image_sampler sampler, const range<dimensions> &rangeRef,
21
                      const property_list &propList = {});
22
23
      /* Available only when: dimensions > 1 */
24
      sampled_image(const void *hostPointer, image_format format,
25
                      image_sampler sampler, const range<dimensions> &rangeRef,
26
                      const range<dimensions -1> &pitch,
27
                      const property_list &propList = {});
28
29
      sampled_image(std::shared_ptr<const void> &hostPointer, image_format format,
30
                      image_sampler sampler, const range<dimensions> &rangeRef,
31
                      const property_list &propList = {});
32
33
      /* Available only when: dimensions > 1 */
34
      sampled_image(std::shared_ptr<const void> &hostPointer, image_format format,
35
                      image_sampler sampler, const range<dimensions> &rangeRef,
36
                      const range<dimensions -1> &pitch,
37
                      const property_list &propList = {});
38
      /* -- common interface members -- */
39
40
41
      /* -- property interface members -- */
42
43
      range<dimensions> get_range() const;
44
45
      /* Available only when: dimensions > 1 */
46
      range<dimensions - 1> get_pitch() const;
```

```
47
48
     size_t get_count() const;
49
50
     size_t get_size() const;
51
52
     template <typename dataT, access::mode accessMode>
53
     accessor<dataT, dimensions, accessMode, access::target::sampled_image>
54
     get_access(handler & commandGroupHandler);
55
56
     template <typename dataT, access::mode accessMode>
57
     accessor<dataT, dimensions, accessMode, access::target::host_sampled_image>
58
     get_access();
59 };
60 } // namespace sycl
```

| Constructor | Description |
|--|---|
| <pre>sampled_image(const void *hostPointer,</pre> | Construct a SYCL sampled_image instance |
| <pre>image_format format,</pre> | with the hostPointer parameter provided. |
| <pre>image_sampler sampler,</pre> | The ownership of this memory is given to |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | the constructed SYCL sampled_image for |
| <pre>const property_list &propList = {})</pre> | the duration of its lifetime. The host ad- |
| | dress is const, so the host accesses must |
| | be read-only. Since, the hostPointer is |
| | const, this image is only initialized with |
| | this memory and there is no write after its |
| | destruction. The element size of the con- |
| | structed SYCL <pre>sampled_image</pre> will be de- |
| | rived from the format parameter. The sam- |
| | pling member function of the constructed |
| | SYCL sampled_image will be derived from |
| | the sampler parameter. The range of |
| | the constructed SYCL sampled_image is |
| | specified by the rangeRef parameter pro- |
| | vided. The pitch of the constructed SYCL |
| | <pre>sampled_image will be the default size de-</pre> |
| | termined by the SYCL runtime. Zero or |
| | more properties can be provided to the con- |
| | structed SYCL sampled_image via an in- |
| | stance of property_list. |
| | Continued on next page |

| Constructor | Description |
|--|---|
| <pre>sampled_image(const void *hostPointer,</pre> | Available only when: dimensions > 1. |
| <pre>image_format format,</pre> | Construct a SYCL <pre>sampled_image</pre> instance |
| <pre>image_sampler sampler,</pre> | with the hostPointer parameter provided. |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | The ownership of this memory is given to |
| <pre>const range<dimensions-1> &pitch,</dimensions-1></pre> | the constructed SYCL sampled_image for |
| <pre>const property_list &propList = {})</pre> | the duration of its lifetime. The host ad- |
| | dress is const, so the host accesses must |
| | be read-only. Since, the hostPointer is |
| | const, this image is only initialized with |
| | this memory and there is no write after de- |
| | struction. The element size of the con- |
| | structed SYCL sampled_image will be de- |
| | rived from the format parameter. The sam- |
| | pling member function of the constructed |
| | SYCL sampled_image will be derived from |
| | the sampler parameter. The range of |
| | the constructed SYCL sampled_image is |
| | specified by the rangeRef parameter pro- |
| | vided. The pitch of the constructed SYCL |
| | <pre>sampled_image will be the pitch param-</pre> |
| | eter provided. Zero or more proper- |
| | ties can be provided to the constructed |
| | SYCL <pre>sampled_image</pre> via an instance of |
| | property_list. |
| <pre>sampled_image(std::shared_ptr<const void="">&</const></pre> | Construct a SYCL <pre>sampled_image</pre> instance |
| hostPointer, | with the hostPointer parameter provided. |
| <pre>image_format format,</pre> | The ownership of this memory is given to |
| <pre>image_sampler sampler,</pre> | the constructed SYCL <pre>sampled_image</pre> for |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | the duration of its lifetime. The host ad- |
| <pre>const property_list &propList = {})</pre> | dress is const, so the host accesses must |
| | be read-only. Since, the hostPointer is |
| | const, this image is only initialized with |
| | this memory and there is no write after its |
| | destruction. The element size of the con- |
| | structed SYCL sampled_image will be de- |
| | rived from the format parameter. The sam- |
| | pling member function of the constructed |
| | SYCL sampled_image will be derived from |
| | the sampler parameter. The range of |
| | the constructed SYCL sampled_image is |
| | specified by the rangeRef parameter pro- |
| | vided. The pitch of the constructed SYCL |
| | <pre>sampled_image will be the default size de-</pre> |
| | termined by the SYCL runtime. Zero or |
| | more properties can be provided to the con- |
| | structed SYCL sampled_image via an in- |
| | stance of property_list. |
| | Continued on next page |

 Table 4.38: Constructors of the sampled_image class template.

| Constructor | Description |
|--|---|
| <pre>sampled_image(std::shared_ptr<const void="">&</const></pre> | Construct a SYCL sampled_image instance |
| hostPointer, | with the hostPointer parameter provided. |
| <pre>image_format format,</pre> | The ownership of this memory is given |
| <pre>image_sampler sampler,</pre> | to the constructed SYCL sampled_image |
| <pre>const range<dimensions> &rangeRef,</dimensions></pre> | for the duration of its lifetime. The |
| <pre>const range<dimensions-1> & pitch,</dimensions-1></pre> | host address is const, so the host accesses |
| <pre>const property_list &propList = {})</pre> | can be read-only. Since, the hostPointer is const, this image is only initialized with this memory and there is no write after its destruction. The element size of the constructed SYCL sampled_image will be derived from the format pa- rameter. The sampling member function of the constructed SYCL sampled_image will be derived from the sampler param- eter. The range of the constructed SYCL sampled_image is specified by the rangeRef parameter provided. The pitch of the con- structed SYCL sampled_image will be the pitch parameter provided. Zero or more properties can be provided to the constructed SYCL sampled_image via an instance of |
| | property_list. |
| | End of table |

| Member function | Description |
|--|---|
| <pre>range<dimensions> get_range()const</dimensions></pre> | Return a range object representing the size |
| | of the image in terms of the number of el- |
| | ements in each dimension as passed to the |
| | constructor. |
| <pre>range<dimensions-1> get_pitch()const</dimensions-1></pre> | Available only when: dimensions > 1. |
| | Return a range object representing the pitch |
| | of the image in bytes. |
| <pre>size_t get_count()const</pre> | Returns the total number of elements in the |
| | <pre>image. Equal to get_range()[0] * *</pre> |
| | <pre>get_range()[dimensions-1].</pre> |
| <pre>size_t get_size()const</pre> | Returns the size of the image storage |
| | in bytes. The number of bytes may |
| | <pre>be greater than get_count()*element size</pre> |
| | due to padding of elements, rows and slices |
| | of the image for efficient access. |
| <pre>template<typename access::mode="" accessmode="" datat,=""></typename></pre> | Returns a valid accessor to the image with |
| <pre>accessor<datat, access::<="" accessmode,="" dimensions,="" pre=""></datat,></pre> | the specified access mode and target. The |
| <pre>target::sampled_image></pre> | only valid types for dataT are int4, uint4, |
| <pre>get_access(handler & commandGroupHandler)</pre> | float4 and half4. |
| | Continued on next page |

Table 4.39: Member functions of the sampled_image class template.

| Member function | Description |
|--|---|
| <pre>template<typename access::mode="" accessmode="" datat,=""></typename></pre> | Returns a valid accessor to the image with |
| <pre>accessor<datat, access::<="" accessmode,="" dimensions,="" pre=""></datat,></pre> | the specified access mode and target. The |
| <pre>target::host_sampled_image></pre> | only valid types for dataT are int4, uint4, |
| <pre>get_access()</pre> | float4 and half4. |
| | End of table |

Table 4.39: Member functions of the sampled_image class template.

4.7.3.3 Image properties

The properties that can be provided when constructing the SYCL unsampled_image and sampled_image classes are describe in Table 4.40.

```
1 namespace sycl {
2
   namespace property {
3
   namespace image {
   class use_host_ptr {
4
5
     public:
        use_host_ptr() = default;
6
7
   };
8
9
   class use_mutex {
10
     public:
11
        use_mutex(std::mutex &mutexRef);
12
13
        std::mutex *get_mutex_ptr() const;
14
   };
15
16
   class context_bound {
17
     public:
        context_bound(context boundContext);
18
19
20
        context get_context() const;
21 };
22 } // namespace image
23 } // namespace property
24 } // namespace sycl
```

| Property | Description |
|--|--|
| <pre>property::image::use_host_ptr</pre> | The use_host_ptr property adds the re- |
| | quirement that the SYCL runtime must not |
| | allocate any memory for the image and in- |
| | stead uses the provided host pointer directly. |
| | This prevents the SYCL runtime from allo- |
| | cating additional temporary storage on the |
| | host. |
| | Continued on next page |

Table 4.40: Properties supported by the SYCL image classes.

| Property | Description |
|---|--|
| <pre>property::image::use_mutex</pre> | The property adds the requirement that the |
| | memory which is owned by the SYCL image |
| | can be shared with the application via a std |
| | ::mutex provided to the property. The std |
| | ::mutex is locked by the runtime whenever |
| | the data is in use and unlocked otherwise. |
| | Data is synchronized with hostData, when |
| | the std::mutex is unlocked by the runtime. |
| <pre>property::image::context_bound</pre> | The context_bound property adds the re- |
| | quirement that the SYCL image can only be |
| | associated with a single SYCL context that |
| | is provided to the property. |
| | End of table |

Table 4.40: Properties supported by the SYCL image classes.

The constructors and member functions of the image property classes are listed in Tables 4.41 and 4.42

| Constructor | Description |
|---|--|
| <pre>property::image::use_host_ptr::use_host_ptr()</pre> | Constructs a SYCL use_host_ptr property |
| | instance. |
| <pre>property::image::use_mutex::use_mutex(std::mutex &</pre> | Constructs a SYCL use_mutex property in- |
| <pre>mutexRef)</pre> | stance with a reference to mutexRef param- |
| | eter provided. |
| <pre>property::image::context_bound::context_bound(</pre> | Constructs a SYCL context_bound prop- |
| <pre>context boundContext)</pre> | erty instance with a copy of a SYCL |
| | context. |
| | End of table |

Table 4.41: Constructors of the image property classes.

| Member function | Description |
|--|---------------------------------------|
| <pre>std::mutex *property::image::use_mutex::</pre> | Returns the std::mutex which was |
| <pre>get_mutex_ptr()const</pre> | specified when constructing this SYCL |
| | use_mutex property. |
| <pre>context property::image::context_bound::get_context</pre> | Returns the context which was spec- |
| ()const | ified when constructing this SYCL |
| | context_bound property. |
| | End of table |

Table 4.42: Member functions of the image property classes.

4.7.3.4 Image synchronization rules

The rules are similar to those described in Section 4.7.2.3.

For the lifetime of the image object, the associated host memory must be left available to the SYCL runtime and

the contents of the associated host memory is unspecified until the image object is destroyed. If an image object value is copied, then only a reference to the underlying image object is copied. The underlying image object is reference-counted. Only after all image value references to the underlying image object have been destroyed is the actual image object itself destroyed.

If an image object is constructed with associated host memory, then its destructor blocks until all operations in all SYCL queues on that image object have completed. Any modifications to the image data will be copied back, if necessary, to the associated host memory. Any errors occurring during destruction are reported to any associated context's asynchronous error handler. If an image object is constructed with a storage object, then the storage object defines what synchronization or copying behavior occurs on image object destruction.

4.7.4 Sharing host memory with the SYCL data management classes

In order to allow the SYCL runtime to do memory management and allow for data dependencies, there are two classes defined, buffer and image. The default behavior for them is that a "raw" pointer is given during the construction of the data management class, with full ownership to use it until the destruction of the SYCL object.

In this section we go in greater detail on sharing or explicitly not sharing host memory with the SYCL data classes, and we will use the buffer class as an example. The same rules will apply to images as well.

4.7.4.1 Default behavior

When using a SYCL buffer, the ownership of the pointer passed to the constructor of the class is, by default, passed to SYCL runtime, and that pointer cannot be used on the host side until the buffer or image is destroyed. A SYCL application can use memory managed by a SYCL buffer within the buffer scope by using a host accessor as defined in 4.7.6. However, there is no guarantee that the host accessor synchronizes with the original host address used in its constructor.

The pointer passed in is the one used to copy data back to the host, if needed, before buffer destruction. The memory pointed by host pointer will not be de-allocated by the runtime, and the data is copied back from the device if there is a need for it.

4.7.4.2 SYCL ownership of the host memory

In the case where there is host memory to be used for initialization of data but there is no intention of using that host memory after the buffer is destroyed, then the buffer can take full ownership of that host memory.

When a buffer owns the host pointer there is no copy back, by default. In this situation, the SYCL application may pass a unique pointer to the host data, which will be then used by the runtime internally to initialize the data in the device.

For example, the following could be used:

```
1 {
2 auto ptr = std::make_unique<int>(-1234);
3 buffer<int, 1> b { std::move(ptr), range { 1 } };
4 // ptr is not valid anymore.
5 // There is nowhere to copy data back
6 }
```

However, optionally the **buffer::set_final_data()** can be set to a **std::weak_ptr** to enable copying data back, to another host memory address that is going to be valid after buffer construction.

```
1 {
2 auto ptr = std::make_unique<int>(-42);
3 buffer<int, 1> b { std::move(ptr), range { 1 } };
4 // ptr is not valid anymore.
5 // There is nowhere to copy data back.
6 // To get copy back, a location can be specified:
7 b.set_final_data(std::weak_ptr<int> { .... })
8 }
```

4.7.4.3 Shared SYCL ownership of the host memory

When an instance of std::shared_ptr is passed to the buffer constructor, then the buffer object and the developer's application share the memory region. If the shared pointer is still used on the application's side then the data will be copied back from the buffer or image and will be available to the application after the buffer or image is destroyed.

If the shared_ptr is not empty, the contents of the referenced memory are used to initialize the buffer. If the shared_ptr is empty, then the buffer is created with uninitialized memory.

When the buffer is destroyed and the data have potentially been updated, if the number of copies of the shared pointer outside the runtime is 0, there is no user-side shared pointer to read the data. Therefore the data is not copied out, and the buffer destructor does not need to wait for the data processes to be finished from OpenCL, as the outcome is not needed on the application's side.

This behavior can be overridden using the set_final_data() member function of the buffer class, which will by any means force the buffer destructor to wait until the data is copied to wherever the set_final_data() member function has put the data (or not wait nor copy if set final data is nullptr).

```
1 {
2
     std::shared_ptr<int> ptr { data };
3
     {
4
       buffer<int, 1> b { ptr, range<2>{ 10, 10 } };
5
       // update the data
6
       [...]
7
     } // Data is copied back because there is an user side shared_ptr
8
  }
1
   {
2
     std::shared_ptr<int> ptr { data };
3
     {
       buffer<int, 1> b { ptr, range<2>{ 10, 10 } };
4
5
       // update the data
6
       [...]
7
       ptr.reset();
     } // Data is not copied back, there is no user side shared_ptr.
8
9
  }
```

4.7.5 Synchronization primitives

When the user wants to use the **buffer** simultaneously in the SYCL runtime and their own code (e.g. a multithreaded mechanism) and wants to use manual synchronization without host accessors, a std::mutex can be passed to the **buffer** constructor via the right **property**.

The runtime promises to lock the mutex whenever the data is in use and unlock it when it no longer needs it.

```
1 {
2
      std::mutex m;
      auto shD = std::make_shared<int>(42)
3
      sycl::buffer b { shD, { sycl::property::buffer::use_mutex { m } } };
4
5
        std::lock_guard lck { m };
6
 7
        // User accesses the data
8
        do_something(shD);
9
        /* m is unlocked when lck goes out of scope, by normal end of this
10
           block but also if an exception is thrown for example */
11
      }
12 }
```

When the runtime releases the mutex the user is guaranteed that the data was copied back on the shared pointer — unless the final data destination has been changed using the member function set_final_data().

4.7.6 Accessors

Accessors provide access to the data managed by a buffer or image, or to local memory allocated by the runtime. Accessors allow users to define **requirements** to memory objects (see Section 3.7.1). Note that construction of an accessor is what defines a memory object requirement, and these requirements are independent of whether there are any uses of an accessor.

The SYCL accessor class template takes the following template parameters:

- A typename specifying the data type that the accessor is providing access to.
- An integer specifying the dimensionality of the accessor.
- A value of access_mode specifying the mode of access the accessor is providing.
- A value of target specifying the target of access the accessor is providing.

The SYCL host_accessor class template takes the following template parameters:

- A typename specifying the data type that the host_accessor is providing access to.
- An integer specifying the dimensionality of the accessor.
- A value of access_mode specifying the mode of access the host_accessor is providing.

The parameters described above determine the data an accessor provides access to and the way in which that access is provided. This separation allows a SYCL runtime implementation to choose an efficient way to provide access to the data within an execution schedule.

Because of this the interfaces of the accessor and the host_accessor will be different depending on the possible combinations of those parameters.

There are four categories of accessor; buffer device accessors (see Section 4.7.6.9), buffer host accessors (see Section 4.7.6.10), local accessors (see Section 4.7.6.11) and image accessors (see Section 4.7.6.12).

4.7.6.1 Access targets

The access target of an accessor specifies what the accessor is providing access to.

The target enumeration, shown in Table 4.43, describes the potential targets of an accessor.

```
namespace sycl {
 1
    enum class target {
2
3
     global_buffer,
4
      constant_buffer,
5
      local,
     unsampled_image,
6
7
     sampled_image,
8
     host_buffer,
9
     host_unsampled_image,
10
     host_sampled_image
11 };
12
13 namespace access {
14
    using sycl::target;
15 } // namespace access
16 } // namespace sycl
```

| target | Description |
|---|---|
| <pre>target::global_buffer</pre> | Access buffer via global memory. |
| <pre>target::constant_buffer</pre> | Access buffer via constant memory. |
| target::local | Access work-group local memory. |
| target::unsampled_image | Access an unsampled_image. |
| target::sampled_image | Access a sampled_image. |
| <pre>target::host_buffer</pre> | Access a buffer immediately in host code. |
| <pre>target::host_unsampled_image</pre> | Access an unsampled_image immediately in |
| | host code. |
| <pre>target::host_sampled_image</pre> | Access a sampled_image immediately in |
| | host code. |
| | End of table |

Table 4.43: Enumeration of access modes available to accessors.

4.7.6.2 Access modes

The access mode of an **accessor** specifies the kind of access that is being provided. This information is used by the runtime to ensure that any data dependencies are resolved by enqueuing any data transfers before or after the execution of a kernel. If a user wants to modify only certain parts of a buffer, preserving other parts of the buffer, then the user should specify the exact sub-range of modification of the buffer.

The access_mode enumeration, shown in Table 4.44, describes the potential modes of an accessor.

```
1 namespace sycl {
   enum class access_mode {
 2
3
     read.
4
     write,
 5
     read_write,
6
      discard_write,
                         // Deprecated in SYCL 2020
7
      discard_read_write, // Deprecated in SYCL 2020
8
      atomic
                         // Deprecated in SYCL 2020
9
   };
10
11 namespace access {
12
    using sycl::access_mode;
13 }
14 } // namespace sycl
```

| access_mode | Description |
|--|--|
| <pre>access_mode::read</pre> | Read-only access. |
| <pre>access_mode::write</pre> | Write-only access. Previous contents not |
| | discarded. |
| <pre>access_mode::read_write</pre> | Read and write access. |
| <pre>access_mode::discard_write</pre> | Deprecated in SYCL 2020. Write-only ac- |
| | cess. Previous contents discarded. |
| <pre>access_mode::discard_read_write</pre> | Deprecated in SYCL 2020. Read and write |
| | access. Previous contents discarded. |
| <pre>access_mode::atomic</pre> | Read and write atomic access. Deprecated |
| | in SYCL 2020. |
| | End of table |

Table 4.44: Enumeration of access modes available to accessors.

4.7.6.3 Access tags

The access mode and access target can be specified via passing tag types to an accessor or host_accessor constructor. This type is used to deduce template arguments of a class.

The Table 4.45, describes the potential tag types and values modes.

SYCL implementations shall provide sufficient list of deduction guides for all template arguments to be deduced for all accessor constructors except for the default constructor and dimension = 0. If accessor template arguments are specified by user, but constructor is called with non-matching tag, the SYCL implementation must emit a compilation error. If a const data type is specified as an accessor template argument and constructor is called with write_only tag, the SYCL implementation must emit a compilation error.

```
1 namespace sycl {
2 template <access_mode>
3 struct mode_tag_t {
4    explicit mode_tag_t() = default;
5 };
6
```
```
read_only{};
 7
   inline constexpr mode_tag_t<access_mode::read>
8
    inline constexpr mode tag t<access mode::read write> read write{}:
9
    inline constexpr mode_tag_t<access_mode::write>
                                                         write_only{};
10
11
   template <access_mode, target>
12
    struct mode_target_tag_t {
13
     explicit mode_target_tag_t() = default;
14
   };
15
16 inline constexpr mode_target_tag_t<access_mode::read, target::constant_buffer> read_constant{};
17 } // namespace sycl
```

| Tag type | Tag value | Access modes | Accessor target |
|------------------------------|---------------|------------------------------------|-------------------------------------|
| <pre>mode_tag_t</pre> | read_write | <pre>access_mode::read_write</pre> | default |
| <pre>mode_tag_t</pre> | read_only | <pre>access_mode::read</pre> | default |
| <pre>mode_tag_t</pre> | write_only | <pre>access_mode::write</pre> | default |
| <pre>mode_target_tag_t</pre> | read_constant | access_mode::read | <pre>target:: constant_buffer</pre> |

Table 4.45: Enumeration of access tags available to accessors.

4.7.6.4 Device and host accessors

There are two types that SYCL applications can use to access data: accessor and host_accessor. Some forms of accessor provide access to data from device kernels while other forms provide immediate access to data from the host. All forms of host_accessor provide access to data from the host, though some forms provide immediate access while other forms provide access to host data from host tasks.

If an accessor has the access target target::global_buffer, target::constant_buffer, target::local, target::unsampled_image or target::sampled_image then it is considered a device accessor, and therefore can only be used within a SYCL kernel function and must be associated with a command group. Creating a device accessor is a non-blocking operation which defines a requirement on the device and adds the requirement to the queue.

Some forms of host_accessor provide immediate access to data from the host, as does accessor when used with access target target::host_unsampled_image, target::host_sampled_image or target::host_buffer. These accessors can only be used from the host and their constructors block until their data is available on the host.

4.7.6.5 Placeholder accessor

Certain accessor types are allowed to be *placeholder* accessors. A *placeholder* accessor defines an accessor instance that is not bound to a specific command group. The accessor defines only the type of the accessor (target memory, access mode, base type, ...). When associated with a command group using the handler::require() function, it defines a **requirement** for the command group. The same placeholder accessor can be required by multiple command groups.

Only the following access targets are allowed in placeholder accessors:

- target::global_buffer
- target::constant_buffer

The enum class placeholder shown below can be used as an accessor template parameter isPlaceholder. This accessor template parameter has been deprecated in SYCL 2020 and placeholder semantics apply regardless of the value isPlaceholder.

```
1 namespace sycl {
2 enum class placeholder { // Deprecated in SYCL 2020
3 false_t,
4 true_t,
5 };
6 } // namespace sycl
```

4.7.6.6 Accessor declaration

The declaration for the accessor and the host_accessor classes is provided in Listing 4.1.

```
1 namespace sycl {
    template <typename dataT,</pre>
 2
 3
              int dimensions = 1,
 4
              access_mode accessmode =
 5
                (std::is_const_v<dataT> ? access_mode::read
 6
                                        : access_mode::read_write),
 7
              target accessTarget = target::global_buffer,
 8
              access::placeholder isPlaceholder = access::placeholder::false_t // Deprecated in SYCL
                  2020
 9 >
10 class accessor;
11
12 template <typename dataT,
13
              int dimensions = 1,
              access_mode accessmode =
14
15
                (std::is_const_v<dataT> ? access_mode::read
16
                                        : access_mode::read_write)
17 >
18 class host_accessor;
19 } // namespace sycl
```

Listing 4.1: Accessor declaration.

4.7.6.7 Constness of the accessor data type

An accessor or host_accessor can be constructed with the underlying data type being const, resulting in an accessor of const dataT. Having an accessor of const dataT is semantically equivalent to having an accessor of access_mode::read: they are both read-only accessors.

Some access modes contradict the constness of the data. An underlying data type const dataT is only valid with the following access modes:

- access_mode::read
- access_mode::read_write

Even when the access mode for an accessor is access_mode::read_write, adding const to dataT makes it a read-only accessor. This ensures, among other things, that the read-only accessor will not trigger a copy-back that

access_mode::read_write normally would require.

Using const dataT makes the accessor read-only by default, as shown in the following example:

```
1 accessor<const int> acc;
2 static_assert(std::is_same_v<
3 decltype(acc),
4 accessor<const int, 1, access_mode::read, target::global_buffer>
5 >);
```

4.7.6.8 Implicit accessor conversions

It is valid to implicitly convert a writable accessor to a read-only one by doing at least one of the following:

- Converting data type from non-const dataT to const dataT
- Converting access mode from access_mode::read_write to access_mode::read

Because of the semantic equivalence defined in 4.7.6.7, the following accessor types can be implicitly converted to one another:

- accessor<dataT, dimensions, access_mode::read, accessTarget>
- accessor<const dataT, dimensions, access_mode::read, accessTarget>
- accessor<const dataT, dimensions, access_mode::read_write, accessTarget>

And the following host_accessor types can be implicitly converted to one another:

- host_accessor<dataT, dimensions, access_mode::read, accessTarget>
- host_accessor<const dataT, dimensions, access_mode::read, accessTarget>
- host_accessor<const dataT, dimensions, access_mode::read_write, accessTarget>

4.7.6.9 Device buffer accessor

A device buffer accessor provides access to a SYCL <u>buffer</u> instance on a device. A SYCL <u>accessor</u> is considered a device buffer accessor if it has the access target <u>target</u>::global_buffer, or <u>target</u>::constant_buffer.

A device buffer accessor can provide access to memory managed by a SYCL buffer class via either global memory or constant memory, corresponding to the access targets target::global_buffer and target:: constant_buffer respectively. A device buffer accessor accessing a SYCL buffer via constant memory is restricted by the available constant memory available on the SYCL device being executed on.

The data type of an accessor must match that of the SYCL buffer which it is accessing.

The dimensionality of a buffer accessor must match that of the SYCL buffer which it is accessing, with the exception of 0 in which case the dimensionality of the SYCL buffer must be 1.

There are three ways a SYCL accessor can provide access to the elements of a SYCL buffer. Firstly by passing a SYCL id instance of the same dimensionality as the SYCL accessor subscript operator. Secondly by passing a single size_t value to multiple consecutive subscript operators (one for each dimension of the SYCL accessor, for example acc[id0][id1][id2]). Finally, in the case of the SYCL accessor being 0 dimensions, by triggering

the implicit conversion operator. Whenever a multi-dimensional index is passed to a SYCL accessor the linear index is calculated based on the index {id0, id1, id2} provided and the range of the SYCL accessor {r0, r1, r2} according to row-major ordering as follows:

$$id2 + (id1 \cdot r2) + (id0 \cdot r2 \cdot r1)$$
 (4.3)

An accessor can optionally provide access to a sub range of a SYCL buffer by providing a range and offset on construction. In this case the SYCL runtime will only guarantee the latest copy of the data is available in that given range and any modifications outside that range are considered undefined behavior. This allows the SYCL runtime to perform optimizations such as reducing copies between devices. The indexing performed when a SYCL accessor provides access to the elements of a SYCL buffer is unaffected, i.e, the accessor will continue to index from $\{0, 0, 0\}$. This allows the offset to be provided either manually or via the parallel_for as shown in Listing 4.2.

```
1
        myQueue.submit([&](handler &cgh) {
2
          auto singleRange = range<3>(8, 16, 16);
3
          auto offset = id<3>(8, 0, 0);
          // We define the subset of the accessor we require for the kernel
4
5
          accessor ptr(syclBuffer, cgh, singleRange, offset);
          // We offset the kernel by the same value to match indexes
6
7
          cgh.parallel_for(singleRange, offset, [=](item<3> itemID) {
            ptr[itemID.get_linear_id()] = 2;
8
9
          });
10
        });
```

Listing 4.2: Using an accessor to access a sub-range.

An accessor with access target target::global_buffer can optionally provide atomic access to a SYCL buffer, using the access_mode::atomic, in which case all operators which return an element of the SYCL buffer return an instance of the deprecated cl::sycl::atomic class. This functionality is provided for backwards compatibility and will be removed in a future version of SYCL.

A device buffer accessor meets the C++ requirement of ContiguousContainer and ReversibleContainer. The exception to this is that the device buffer accessor destructor doesn't destroy any elements or free the memory, because an accessor doesn't own the underlying data. The iterator for the container interface is the same pointer type as obtained by calling get_multi_ptr<access::decorated::no>(). For multidimensional accessors the iterator linearizes the data according to 4.3.

The full list of capabilities that device buffer accessors can support is described in 4.46.

| Access target | Access modes | Data types | Dimensionalities |
|-----------------|---------------------------------------|---|-------------------------------------|
| global_buffer | read write read_write atomic | The data type of the SYCL buffer being accessed. | Between 0 and 3 (inclusive). |
| constant_buffer | read | The data type of the SYCL buffer be- ing accessed. | Between 0 and 3 (inclusive). |

Table 4.46: Description of all the device **buffer accessor** capabilities.

4.7.6.9.1 Device buffer accessor interface

A synopsis of the SYCL accessor class template buffer specialization is provided below. The member types for this accessor specialization are listed in Tables 4.47. The constructors for this accessor specialization are listed in Tables 4.48. The member functions for this accessor specialization are listed in Tables 4.49. The additional common special member functions and common member functions are listed in 4.5.3 in Tables 4.1 and 4.2, respectively. For valid implicit conversions between accessor types please refer to 4.7.6.8. Additionally, accessors of the same type must be equality comparable not only on the host application, but also within SYCL kernel functions.

```
1 namespace sycl {
2
   template <typename dataT,</pre>
3
              int dimensions,
4
              access::mode accessmode,
5
              access::target accessTarget,
              access::placeholder isPlaceholder>
6
7
   class accessor {
8
     public:
9
      template <access::decorated IsDecorated>
      using accessor_ptr = // Corresponds to the target address space,
10
11
          __pointer_class__; // is pointer-to-const
                             // when (accessmode == access::mode::read);
12
13
     using value_type = // const dataT when (accessmode == access::mode::read),
          __value_type__; // dataT otherwise
14
15
     using reference =
                             // const dataT& when (accessmode == access::mode::read),
          __reference_type__; // dataT& otherwise
16
17
      using const_reference = const dataT &;
18
      using iterator =
                          // Corresponds to the target address space,
          __pointer_type__; // is pointer-to-const
19
20
                           // when (accessmode == access::mode::read)
21
      using const_iterator =
          __pointer_to_const_type__; // Corresponds to the target address space
22
23
      using reverse_iterator = std::reverse_iterator<iterator>;
24
      using const_reverse_iterator = std::reverse_iterator<const_iterator>;
25
      using difference_type =
26
          typename std::iterator_traits<iterator>::difference_type;
27
      using size_type = size_t;
28
29
      accessor();
30
31
      /* Available only when: (dimensions == 0) */
32
      template <typename AllocatorT>
33
      accessor(buffer<dataT, 1, AllocatorT> &bufferRef,
34
               const property_list &propList = {});
35
      /* Available only when: (dimensions == 0) */
36
37
      template <typename AllocatorT>
38
      accessor(buffer<dataT, 1, AllocatorT> &bufferRef,
39
               handler &commandGroupHandlerRef, const property_list &propList = {});
40
41
      /* Available only when: (dimensions > 0) */
42
      template <typename AllocatorT>
43
      accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
44
               const property_list &propList = {});
```

```
45
46
      /* Available only when: (dimensions > 0) */
47
      template <typename AllocatorT, typename TagT>
48
      accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef, TagT tag,
49
               const property_list &propList = {});
50
51
      /* Available only when: (dimensions > 0) */
52
      template <typename AllocatorT>
53
      accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
54
               handler &commandGroupHandlerRef, const property_list &propList = {});
55
      /* Available only when: (dimensions > 0) */
56
57
      template <typename AllocatorT, typename TagT>
58
      accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
59
               handler &commandGroupHandlerRef, TagT tag,
60
               const property_list &propList = {});
61
62
      /* Available only when: (dimensions > 0) */
63
      template <typename AllocatorT>
64
      accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
65
               range<dimensions> accessRange, const property_list &propList = {});
66
67
      /* Available only when: (dimensions > 0) */
68
      template <typename AllocatorT, typename TagT>
69
      accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
70
               range<dimensions> accessRange, TagT tag,
71
               const property_list &propList = {});
72
73
      /* Available only when: (dimensions > 0) */
74
      template <typename AllocatorT>
75
      accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
76
               range<dimensions> accessRange, id<dimensions> accessOffset,
77
               const property_list &propList = {});
78
79
      /* Available only when: (dimensions > 0) */
80
      template <typename AllocatorT, typename TagT>
      accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
81
82
               range<dimensions> accessRange, id<dimensions> accessOffset,
83
               TagT tag, const property_list &propList = {});
84
85
      /* Available only when: (dimensions > 0) */
86
      template <typename AllocatorT>
87
      accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
88
               handler &commandGroupHandlerRef, range<dimensions> accessRange,
89
               const property_list &propList = {});
90
91
      /* Available only when: (dimensions > 0) */
92
      template <typename AllocatorT, typename TagT>
93
      accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
94
               handler &commandGroupHandlerRef, range<dimensions> accessRange,
95
               TagT tag, const property_list &propList = {});
96
97
      /* Available only when: (dimensions > 0) */
98
      template <typename AllocatorT>
99
      accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
```

```
100
                handler &commandGroupHandlerRef, range<dimensions> accessRange,
101
                id<dimensions> accessOffset, const property_list &propList = {});
102
103
       /* Available only when: (dimensions > 0) */
104
       template <typename AllocatorT, typename TagT>
105
       accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
106
                handler &commandGroupHandlerRef, range<dimensions> accessRange,
107
                id<dimensions> accessOffset, TagT tag,
108
                const property_list &propList = {});
109
       /* -- common interface members -- */
110
111
       void swap(accessor &other);
112
113
114
       bool is_placeholder() const;
115
116
       size_type byte_size() const noexcept;
117
118
       size_type size() const noexcept;
119
120
       size_type max_size() const noexcept;
121
122
       size_type get_count() const noexcept;
123
124
       bool empty() const noexcept;
125
126
       /* Available only when: dimensions > 0 */
127
       range<dimensions> get_range() const;
128
129
       /* Available only when: dimensions > 0 */
130
       id<dimensions> get_offset() const;
131
132
       /* Available only when: (dimensions == 0) */
133
       operator reference() const;
134
135
       /* Available only when: (dimensions > 0) */
136
       reference operator[](id<dimensions> index) const;
137
138
       /* Deprecated in SYCL 2020
139
       Available only when: accessMode == access::mode::atomic && dimensions == 0 */
140
       operator cl::sycl::atomic<dataT, access::address_space::global_space> () const;
141
142
       /* Deprecated in SYCL 2020
143
       Available only when: accessMode == access::mode::atomic && dimensions > 0 */
144
       cl::sycl::atomic<dataT, access::address_space::global_space> operator[](
145
         id<dimensions> index) const;
146
147
       /* Available only when: dimensions > 1 */
148
       __unspecified__ &operator[](size_t index) const;
149
150
       std::add_pointer_t<value_type> get_pointer() const noexcept;
151
152
       template <access::decorated IsDecorated>
153
       accessor_ptr<IsDecorated> get_multi_ptr() const noexcept;
154
```

```
155
       iterator data() const noexcept;
156
157
       iterator begin() const noexcept;
158
159
       iterator end() const noexcept;
160
       const_iterator cbegin() const noexcept;
161
162
163
       const_iterator cend() const noexcept;
164 };
165
166 } // namespace sycl
```

Listing 4.3: Device accessor class for buffers.

| Member types | Description |
|--|--|
| value_type | If (accessmode == access_mode::read), |
| | equal to const dataT. In other cases equal |
| | to dataT. |
| <pre>template <access::decorated isdecorated=""></access::decorated></pre> | <pre>If (accessTarget == access::target::</pre> |
| accessor_ptr | <pre>global_buffer): multi_ptr<value_type,< pre=""></value_type,<></pre> |
| | <pre>access::address_space::global_space,</pre> |
| | IsDecorated>. |
| | If (accessTarget == access::target |
| | ::constant_buffer): multi_ptr< |
| | <pre>value_type, access::address_space</pre> |
| | ::constant_space, IsDecorated>. |
| reference | If (accessmode == access_mode::read), |
| | equal to const dataT&. In other cases equal |
| | to dataT&. |
| const_reference | const dataT& |
| iterator | If (accessTarget == access::target |
| | ::global_buffer): raw_global_ptr< |
| | value_type>. |
| | If (accessTarget == access::target:: |
| | <pre>constant_buffer): raw_constant_ptr<</pre> |
| | value_type>. |
| const_iterator | <pre>If (accessTarget == access::target::</pre> |
| | <pre>global_buffer): raw_global_ptr<const< pre=""></const<></pre> |
| | value_type>. |
| | If (accessTarget == access::target:: |
| | <pre>constant_buffer): raw_constant_ptr<</pre> |
| | <pre>const value_type>.</pre> |
| reverse_iterator | Iterator adaptor that reverses the direction of |
| | iterator. |
| const_reverse_iterator | Iterator adaptor that reverses the direction of |
| | const_iterator. |
| difference_type | <pre>typename std::iterator_traits<</pre> |
| | <pre>iterator>::difference_type</pre> |
| size_type | size_t |
| | End of table |

Table 4.47: Member types of the accessor class template buffer specialization.

| Constructor | Description |
|---|--|
| accessor() | Constructs an empty accessor. Fulfills the |
| | following post-conditions: |
| | <pre>• (empty()== true)</pre> |
| | • All size queries return 0. |
| | • The only iterator that can be obtained |
| | is nullptr. |
| | • Trying to access the underlying mem- |
| | ory is undefined behavior. |
| | A default constructed placeholder accessor |
| | can be passed to a SYCL kernel, but it is not |
| | valid to register it with the command group |
| | handler. |
| <pre>accessor(buffer<datat, 1,="" allocatort=""> &bufferRef,</datat,></pre> | Available only when: (dimensions $== 0$). |
| <pre>const property_list &propList = {})</pre> | Constructs a placeholder accessor instance |
| | for accessing the first element of a SYCL |
| | <pre>buffer. The optional property_list pro-</pre> |
| | vides properties for the constructed SYCL |
| | accessor object. |
| <pre>accessor(buffer<datat, 1,="" allocatort=""> &bufferRef,</datat,></pre> | Available only when: (dimensions $== 0$). |
| handler &commandGroupHandlerRef, | Constructs a SYCL accessor instance |
| <pre>const property_list &propList = {})</pre> | for accessing the first element of a SYCL |
| | buffer within a SYCL kernel function |
| | on the SYCL queue associated with |
| | commandGroupHandlerRef. The optional |
| | property_list provides properties for the |
| | constructed SYCL accessor object. |
| accessor(buffer <datat, allocatort="" dimensions,=""> &</datat,> | Available only when: (dimensions > 0). |
| bufferKef, | Constructs a placeholder accessor for ac- |
| <pre>const property_list &propList = {})</pre> | cessing a SYCL buffer. The optional |
| | property_list provides properties for the |
| accordent for data dimensions Allocator The | Available only when (dimensions) () |
| accessor(buffer <datal, allocatori="" dimensions,=""> &</datal,> | Available only when: (dimensions > 0). |
| The second property list (proplict - ()) | constructs a placeholder accessor for ac- |
| Tagi tag, const property_fist aproprist = {}) | deduce template arguments of the accessor |
| | as described in Section 4.7.6.3. The optional |
| | property list provides properties for the |
| | constructed SYCL accessor object |
| accessor(buffer <datat_dimensions_allocatort> &</datat_dimensions_allocatort> | Available only when: $(dimensions > 0)$ |
| bufferRef. | Constructs a SYCL accessor instance for |
| handler &commandGroupHandlerRef. | accessing a SYCL buffer within a SYCL |
| <pre>const property list &propList = {})</pre> | kernel function on the SYCL queue asso- |
| | ciated with commandGroupHandlerRef. The |
| | optional property_list provides properties |
| | for the constructed SYCL accessor object. |
| | Continued on next page |

Table 4.48: Constructors of the accessor class template buffer specialization.

| Constructor | Description |
|--|---|
| <pre>accessor(buffer<datat, allocatort="" dimensions,=""> &</datat,></pre> | Available only when: (dimensions > 0). |
| bufferRef, | Constructs a SYCL accessor instance for |
| handler &commandGroupHandlerRef, | accessing a SYCL buffer within a SYCL |
| <pre>TagT tag, const property_list &propList = {})</pre> | kernel function on the SYCL queue asso- |
| | ciated with commandGroupHandlerRef. The |
| | tag is used to deduce template arguments of |
| | the accessor as described in Section 4.7.6.3. |
| | The optional property_list provides prop- |
| | erties for the constructed SYCL accessor |
| | object. |
| <pre>accessor(buffer<datat, allocatort="" dimensions,=""> &</datat,></pre> | Available only when: (dimensions > 0). |
| bufferRef, | Constructs a placeholder accessor for ac- |
| range <dimensions> accessRange,</dimensions> | cessing a range of a SYCL buffer. The |
| <pre>const property_list &propList = {})</pre> | optional property_list provides properties |
| | for the constructed SYCL accessor object. |
| accessor(buffer <datal, allocatori="" dimensions,=""> &</datal,> | Available only when: $(dimensions > 0)$. |
| buiierRei, | constructs a placeholder accessor for ac- |
| range <dimensions> accessRange,</dimensions> | cessing a range of a STCL buffer. The |
| Tagi tag, const property_fist aprophist = {}) | the accessor as described in Section 4.7.6.3 |
| | The optional property list provides prop |
| | erties for the constructed SVCL accessor |
| | object |
| accessor(buffer <datat, allocatort="" dimensions,=""> &</datat,> | Available only when: $(dimensions > 0)$. |
| bufferRef. | Constructs a placeholder accessor for ac- |
| <pre>range<dimensions> accessRange,</dimensions></pre> | cessing a range of a SYCL buffer. The |
| <pre>id<dimensions> accessOffset,</dimensions></pre> | optional property_list provides properties |
| <pre>const property_list &propList = {})</pre> | for the constructed SYCL accessor object. |
| <pre>accessor(buffer<datat, allocatort="" dimensions,=""> &</datat,></pre> | Available only when: (dimensions > 0). |
| bufferRef, | Constructs a placeholder accessor for ac- |
| <pre>range<dimensions> accessRange,</dimensions></pre> | cessing a range of a SYCL buffer. The |
| <pre>id<dimensions> accessOffset,</dimensions></pre> | tag is used to deduce template arguments of |
| <pre>TagT tag, const property_list &propList = {})</pre> | the accessor as described in Section 4.7.6.3. |
| | The optional property_list provides prop- |
| | erties for the constructed SYCL accessor |
| | object. |
| <pre>accessor(buffer<datat, allocatort="" dimensions,=""> &</datat,></pre> | Available only when: (dimensions > 0). |
| bufferRef, | Constructs a SYCL accessor instance for |
| handler &commandGroupHandlerRef, | accessing a range of SYCL buffer within |
| <pre>range<dimensions> accessRange,</dimensions></pre> | a SYCL kernel function on the SYCL queue |
| <pre>const property_list &propList = {})</pre> | associated with commandGroupHandlerRef, |
| | specified by accessRange. The optional |
| | property_list provides properties for the |
| | constructed STCL accessor object. |
| | Continued on next page |

Table 4.48: Constructors of the accessor class template buffer specialization.

| Constructor | Description |
|--|--|
| <pre>accessor(buffer<datat, allocatort="" dimensions,=""> &</datat,></pre> | Available only when: (dimensions > 0). |
| bufferRef, | Constructs a SYCL accessor instance for |
| <pre>handler &commandGroupHandlerRef,</pre> | accessing a range of SYCL buffer within |
| <pre>range<dimensions> accessRange,</dimensions></pre> | a SYCL kernel function on the SYCL queue |
| <pre>TagT tag, const property_list &propList = {})</pre> | associated with commandGroupHandlerRef, |
| | specified by accessRange. The tag is |
| | used to deduce template arguments of the |
| | accessor as described in Section 4.7.6.3. |
| | The optional property_list provides prop- |
| | erties for the constructed SYCL accessor |
| | object. |
| <pre>accessor(buffer<datat, allocatort="" dimensions,=""> &</datat,></pre> | Available only when: (dimensions > 0). |
| bufferRef, | Constructs a SYCL accessor in- |
| <pre>handler &commandGroupHandlerRef,</pre> | stance for accessing a range of SYCL |
| <pre>range<dimensions> accessRange,</dimensions></pre> | buffer within a SYCL kernel function |
| <pre>id<dimensions> accessOffset,</dimensions></pre> | on the SYCL queue associated with |
| <pre>const property_list &propList = {})</pre> | commandGroupHandlerRef, specified by |
| | accessRange and accessOffset. The |
| | optional property_list provides properties |
| | for the constructed SYCL accessor object. |
| <pre>accessor(buffer<datat, allocatort="" dimensions,=""> &</datat,></pre> | Available only when: (dimensions > 0). |
| bufferRef, | Constructs a SYCL accessor in- |
| handler &commandGroupHandlerRef, | stance for accessing a range of SYCL |
| <pre>range<dimensions> accessRange,</dimensions></pre> | buffer within a SYCL kernel function |
| <pre>id<dimensions> accessOffset,</dimensions></pre> | on the SYCL queue associated with |
| <pre>TagT tag, const property_list &propList = {})</pre> | commandGroupHandlerRef, specified by |
| | accessRange and accessOffset. The tag is |
| | used to deduce template arguments of the |
| | accessor as described in Section 4.7.6.3. |
| | The optional property_list provides prop- |
| | erties for the constructed SYCL accessor |
| | object. |
| | End of table |

Table 4.48: Constructors of the accessor class template buffer specialization.

| Member function | Description |
|--|--|
| <pre>void swap(accessor &other);</pre> | Swaps the contents of the current accessor |
| | with the contents of other. |
| <pre>bool is_placeholder()const</pre> | Returns true if (accessTarget != target |
| | ::host_buffer) and the accessor has been |
| | constructed without a handler. Otherwise re- |
| | turns false. |
| <pre>size_type byte_size()const noexcept</pre> | Returns the size in bytes of the region of a |
| | SYCL buffer that this SYCL accessor is |
| | accessing. |
| | Continued on next page |

Table 4.49: Member functions of the accessor class template buffer specialization.

| Member function | Description |
|---|--|
| <pre>size_type size()const noexcept</pre> | Returns the number of elements in the re- |
| | gion of a SYCL buffer that this SYCL |
| | accessor is accessing. |
| <pre>size_type max_size()const noexcept</pre> | Returns the maximum number of elements |
| | any accessor of this type would be able to |
| | access. |
| <pre>size_type get_count()const noexcept</pre> | Returns the same as size(). |
| <pre>bool empty()const noexcept</pre> | Returns true iff (size()== 0). |
| <pre>range<dimensions> get_range()const</dimensions></pre> | Available only when: dimensions > 0 . |
| | Returns a range object which represents the |
| | number of elements of dataT per dimension |
| | that this accessor may access. |
| | The range object returned must equal to the |
| | range of the buffer this accessor is associ- |
| | aled with, unless a range was explicitly spec- |
| id dimensiones, not offerst () const | Available only when dimensions > 0 |
| iu <aimensions> get_offset()const</aimensions> | Available only when $d = d = d = d = d = d = d = d = d = d $ |
| | starting point in number of elements of |
| | dataT for the range that this accessor may |
| | access |
| | The id object returned must equal to $id\{0,$ |
| | 0. 0}. unless an offset was explicitly speci- |
| | fied when this accessor was constructed. |
| operator reference()const | Available only when: (dimensions == 0). |
| | Returns a reference to the element stored |
| | within the SYCL buffer this SYCL |
| | accessor is accessing. |
| <pre>reference operator[](id<dimensions> index)const</dimensions></pre> | Available only when: (dimensions > 0). |
| | Returns a reference to the element stored |
| | within the SYCL buffer this SYCL |
| | accessor is accessing at the index specified |
| | by index. |
| operator const dataT &()const | Available only when: accessMode == |
| | <pre>access_mode::read && dimensions == 0)</pre> |
| | |
| | Returns a const reference to the element |
| | stored within the SYCL buffer this SYCL |
| const reference exercise[](id dimensione) index) | Available only when accessMode |
| const_reference operator[](10<01mensions> index) | Available only whell: accessmode == |
| | Returns a const reference to the element \mathcal{R} |
| | stored within the SYCI buffer this SYCI |
| | accessor is accessing at the index specified |
| | hy index |
| | Continued on next page |
| | stored within the SYCL buffer this SYCL accessor is accessing at the index specified by index. Continued on next page |

Table 4.49: Member functions of the accessor class template buffer specialization.

| Member function | Description |
|---|---|
| <pre>operator cl::sycl::atomic<datat,< pre=""></datat,<></pre> | Deprecated in SYCL 2020. |
| <pre>access::address_space::global_space> ()const</pre> | Available only when: accessMode == |
| | <pre>access_mode::atomic && dimensions ==</pre> |
| | 0). |
| | Returns an instance of cl::sycl::atomic |
| | of type dataT providing atomic access to |
| | the element stored within the SYCL buffer |
| | this SYCL accessor is accessing. |
| <pre>cl::sycl::atomic<datat, access::address_space::<="" pre=""></datat,></pre> | Deprecated in SYCL 2020. |
| global_space> | Available only when: accessMode == |
| <pre>operator[](id<dimensions> index)const</dimensions></pre> | <pre>access_mode::atomic && dimensions ></pre> |
| | 0). |
| | Returns an instance of cl::sycl::atomic |
| | of type dataT providing atomic access to the |
| | element stored within the SYCL buffer this |
| | SYCL accessor is accessing at the index |
| | specified by index. |
| unspecified &operator[](size_t index)const | Available only when: dimensions > 1 . |
| | Returns an instance of an undefined inter- |
| | mediate type representing a SYCL accessor |
| | of the same type as this SYCL accessor, |
| | with the dimensionality dimensions-1 and |
| | containing an implicit SYCL id with index |
| | dimensions set to index. The intermedi- |
| | ate type returned must provide all available |
| | subscript operators which take a size_t pa- |
| | rameter defined by the SYCL accessor class |
| | that are appropriate for the type it represents |
| | (including this subscript operator). |
| <pre>std::add_pointer_t<value_type> get_pointer()const</value_type></pre> | Returns a pointer to the memory this SYCL |
| noexcept | accessor memory is accessing. |
| <pre>template <access::decorated isdecorated=""></access::decorated></pre> | Returns a multi_ptr to the memory this |
| accessor_ptr <isdecorated> get_multi_ptr()const</isdecorated> | SYCL accessor memory is accessing. |
| noexcept | |
| iterator data()const noexcept | Returns a pointer to the memory this SYCL |
| | accessor memory is accessing. |
| <pre>iterator begin()const noexcept</pre> | Returns an iterator to the first element of the |
| | memory within the access range. |
| iterator end()const noexcept | Returns an iterator that points past the last |
| | element of the memory within the access |
| | range. |
| <pre>const_iterator cbegin()const noexcept</pre> | Returns a const iterator to the first element |
| | of the memory within the access range. |
| const iterator cend()const noexcept | Returns a const iterator that points past the |
| | last element of the memory within the access |
| | range. |
| | Continued on next page |

Table 4.49: Member functions of the **accessor** class template buffer specialization.

| Member function | Description |
|---|--|
| reverse_iterator rbegin()const noexcept | Returns an iterator adaptor to the last ele- |
| | ment of the memory within the access range. |
| reverse_iterator rend()const noexcept | Returns an iterator adaptor that points before |
| | the first element of the memory within the |
| | access range. |
| <pre>const_reverse_iterator crbegin()const noexcept</pre> | Returns a const iterator adaptor to the last |
| | element of the memory within the access |
| | range. |
| <pre>const_reverse_iterator crend()const noexcept</pre> | Returns a const iterator adaptor that points |
| | before the first element of the memory |
| | within the access range. |
| | End of table |

Table 4.49: Member functions of the accessor class template buffer specialization.

4.7.6.9.2 Device buffer accessor properties

The properties that can be provided when constructing the SYCL accessor class are describe in Table 4.50.

```
1 namespace sycl {
2 namespace property {
3 struct noinit {};
4 } // namespace property
5
6 inline constexpr property::noinit noinit;
7 } // namespace sycl
```

| Property | Description |
|-----------------------------|--|
| <pre>property::noinit</pre> | The noinit property notifies the SYCL run- |
| | time that previous contents of a buffer can be |
| | discarded. |
| | Replaces deprecated discard_write and |
| | discard_read_write access modes. |
| | End of table |

Table 4.50: Properties supported by the SYCL accessor class.

The constructors of the accessor property classes are listed in Table 4.51.

| Constructor | Description |
|---------------------------------------|---------------------------------------|
| <pre>property::noinit::noinit()</pre> | Constructs a SYCL noinit property in- |
| | stance. |
| | End of table |

Table 4.51: Constructors of the accessor property classes.

4.7.6.10 Host buffer accessor

A SYCL host_accessor is a host buffer accessor, which provides access to a SYCL buffer instance on a host.

A host_accessor can be used in either of two ways. If the accessor is used from application scope, it must be constructed without a command group handler object. In this case, the constructor blocks until the requested memory is copied to the host, and the accessor provides immediate access to that memory.

A host_accessor can also be used inside command group scope for a host task (see Section 4.11). In this case, it must be constructed with the command group handler for the host task. The constructor does not block in this case. Instead, the accessor provides a requisite for scheduling the host task.

If the SYCL buffer this SYCL host_accessor is accessing was constructed with the property property::buffer ::use_host_ptr the address of the memory accessed on the host must be the address the SYCL buffer was constructed with, otherwise the SYCL runtime is free to allocate temporary memory to provide access on the host.

The data type of a host buffer accessor must match that of the SYCL buffer which it is accessing.

The dimensionality of a buffer accessor must match that of the SYCL buffer which it is accessing, with the exception of 0 in which case the dimensionality of the SYCL buffer must be 1.

There are three ways a SYCL host_accessor can provide access to the elements of a SYCL buffer. Firstly by passing a SYCL id instance of the same dimensionality as the SYCL host_accessor subscript operator. Secondly by passing a single size_t value to multiple consecutive subscript operators (one for each dimension of the SYCL host_accessor, for example acc[id0][id1][id2]). Finally, in the case of the SYCL host_accessor being 0 dimensions, by triggering the implicit conversion operator. Whenever a multi-dimensional index is passed to a SYCL host_accessor the linear index is calculated based on the index {id0, id1, id2} provided and the range of the SYCL host_accessor {r0, r1, r2} according to row-major ordering as follows:

$$id2 + (id1 \cdot r2) + (id0 \cdot r2 \cdot r1) \tag{4.4}$$

A host buffer accessor can optionally provide access to a sub range of a SYCL buffer by providing a range and offset on construction. In this case the SYCL runtime will only guarantee the latest copy of the data is available in that given range and any modifications outside that range are considered undefined behavior. The indexing performed when a SYCL host_accessor provides access to the elements of a SYCL buffer is unaffected, i.e, the accessor will continue to index from {0,0,0}.

A host buffer accessor meets the C++ requirement of ContiguousContainer and ReversibleContainer. The exception to this is that the device buffer accessor destructor doesn't destroy any elements or free the memory, because a host buffer accessor doesn't own the underlying data. For multidimensional accessors the iterator linearizes the data according to 4.3.

4.7.6.10.1 Host buffer accessor interface

A synopsis of the SYCL host_accessor class template buffer specialization is provided below. The member types for this host_accessor specialization are listed in Tables 4.52. The constructors for this host_accessor specialization are listed in Tables 4.53. The member functions for this host_accessor specialization are listed in Tables 4.54. The additional common special member functions and common member functions are listed in 4.5.3 in Tables 4.1 and 4.2, respectively. For valid implicit conversions between accessor types please refer to 4.7.6.8. Additionally, accessors of the same type must be equality comparable.

```
1 namespace sycl {
 2 template <typename dataT,</pre>
 3
              int dimensions,
 4
              access::mode accessmode>
 5
    class host_accessor {
 6
     public:
 7
      using value_type = // const dataT when (accessmode == access::mode::read),
          __value_type__; // dataT otherwise
 8
 9
      using reference = // const dataT& when (accessmode == access::mode::read),
10
          __reference_type__; // dataT& otherwise
      using const_reference = const dataT &;
11
      using iterator = // const dataT* when (accessmode == access::mode::read),
12
          __pointer_type__; // dataT* otherwise
13
14
      using const_iterator = const dataT *;
15
      using difference_type =
16
          typename std::iterator_traits<iterator>::difference_type;
17
      using size_type = size_t;
18
19
      host_accessor();
20
21
      /* Available only when: (dimensions == 0) */
      template <typename AllocatorT>
22
23
      host_accessor(buffer<dataT, 1, AllocatorT> &bufferRef,
24
                    const property_list &propList = {});
25
26
      /* Available only when: (dimensions == 0) */
27
      template <typename AllocatorT>
28
      host_accessor(buffer<dataT, 1, AllocatorT> &bufferRef,
29
                    handler &commandGroupHandlerRef, const property_list &propList = {});
30
31
      /* Available only when: (dimensions > 0) */
32
      template <typename AllocatorT>
33
      host_accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
34
                    const property_list &propList = {});
35
36
      /* Available only when: (dimensions > 0) */
37
      template <typename AllocatorT, typename TagT>
38
      host_accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef, TagT tag,
39
                    const property_list &propList = {});
40
41
      /* Available only when: (dimensions > 0) */
42
      template <typename AllocatorT>
43
      host_accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
44
                    handler &commandGroupHandlerRef, const property_list &propList = {});
45
46
      /* Available only when: (dimensions > 0) */
47
      template <typename AllocatorT, typename TagT>
48
      host_accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
49
                    handler &commandGroupHandlerRef, TagT tag,
50
                    const property_list &propList = {});
51
      /* Available only when: (dimensions > 0) */
52
53
      template <typename AllocatorT>
54
      host_accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef,
55
                    range<dimensions> accessRange, const property_list &propList = {});
```

56 57 /* Available only when: (dimensions > 0) */ 58 template <typename AllocatorT, typename TagT> 59 host_accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef, 60 range<dimensions> accessRange, TagT tag, 61 const property_list &propList = {}); 62 63 /* Available only when: (dimensions > 0) */ 64 template <typename AllocatorT> 65 host_accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef, 66 range<dimensions> accessRange, id<dimensions> accessOffset, 67 const property_list &propList = {}); 68 69 /* Available only when: (dimensions > 0) */ 70 template <typename AllocatorT, typename TagT> 71 host_accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef, 72 range<dimensions> accessRange, id<dimensions> accessOffset, 73 TagT tag, const property_list &propList = {}); 74 75 /* Available only when: (dimensions > 0) */ 76 template <typename AllocatorT> host_accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef, 77 78 handler &commandGroupHandlerRef, range<dimensions> accessRange, 79 const property_list &propList = {}); 80 /* Available only when: (dimensions > 0) */ 81 82 template <typename AllocatorT, typename TagT> host_accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef, 83 84 handler &commandGroupHandlerRef, range<dimensions> accessRange, 85 TagT tag, const property_list &propList = {}); 86 87 /* Available only when: (dimensions > 0) */ 88 template <typename AllocatorT> 89 host_accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef, 90 handler &commandGroupHandlerRef, range<dimensions> accessRange, 91 id<dimensions> accessOffset, const property_list &propList = {}); 92 93 /* Available only when: (dimensions > 0) */ 94 template <typename AllocatorT, typename TagT> 95 host_accessor(buffer<dataT, dimensions, AllocatorT> &bufferRef, 96 handler &commandGroupHandlerRef, range<dimensions> accessRange, 97 id<dimensions> accessOffset, TagT tag, 98 const property_list &propList = {}); 99 100 /* -- common interface members -- */ 101 102 void swap(host_accessor &other); 103 104 size_type byte_size() const noexcept; 105 106 size_type size() const noexcept; 107 108 size_type max_size() const noexcept; 109 110 bool empty() const noexcept;

```
111
112
       /* Available only when: dimensions > 0 */
113
       range<dimensions> get_range() const;
114
115
       /* Available only when: dimensions > 0 */
       id<dimensions> get_offset() const;
116
117
       /* Available only when: (dimensions == 0) */
118
119
       operator reference() const;
120
121
       /* Available only when: (dimensions > 0) */
122
       reference operator[](id<dimensions> index) const;
123
124
       /* Available only when: dimensions > 1 */
125
       __unspecified__ &operator[](size_t index) const;
126
127
       iterator data() const noexcept;
128
129
       iterator begin() const noexcept;
130
131
       iterator end() const noexcept;
132
133
       const_iterator cbegin() const noexcept;
134
135
       const_iterator cend() const noexcept;
136 };
137 } // namespace sycl
```

Listing 4.4: Host accessor class for buffers.

| Member types | Description |
|-----------------|--|
| value_type | <pre>If (accessmode == access_mode::read),</pre> |
| | equal to const dataT. In other cases equal |
| | to dataT. |
| reference | <pre>If (accessmode == access_mode::read),</pre> |
| | equal to const dataT&. In other cases equal |
| | to dataT&. |
| const_reference | const dataT& |
| iterator | <pre>If (accessmode == access_mode::read),</pre> |
| | equal to const dataT*. In other cases equal |
| | to dataT*. |
| const_iterator | const dataT* |
| difference_type | <pre>typename std::iterator_traits<</pre> |
| | iterator>::difference_type |
| size_type | size_t |
| | End of table |

Table 4.52: Member types of the host_accessor class template .

| Constructor | Description |
|---|--|
| host_accessor() | Constructs an empty accessor. Fulfills the |
| | following post-conditions: |
| | <pre>• (empty()== true)</pre> |
| | All size queries return 0. |
| | • The only iterator that can be obtained |
| | is nullptr. |
| | • Trying to access the underlying mem- |
| | ory is undefined behavior. |
| | |
| <pre>host_accessor(buffer<datat, 1,="" allocatort=""> &</datat,></pre> | Available only when: (dimensions $== 0$). |
| bufferRef, | Constructs a host_accessor instance for ac- |
| <pre>const property_list &propList = {})</pre> | cessing the first element of a SYCL buffer |
| | immediately on the host. The optional |
| | <pre>property_list provides properties for the</pre> |
| | constructed SYCL host_accessor object. |
| <pre>host_accessor(buffer<datat, 1,="" allocatort=""> &</datat,></pre> | Available only when: (dimensions $== 0$). |
| bufferRef, | Constructs a host_accessor for accessing |
| <pre>handler &commandGroupHandlerRef,</pre> | the first element of a SYCL buffer inside a |
| <pre>const property_list &propList = {})</pre> | host task. The optional property_list pro- |
| | vides properties for the constructed SYCL |
| | host_accessor object. |
| <pre>host_accessor(buffer<datat, allocatort="" dimensions,=""></datat,></pre> | Available only when: (dimensions > 0). |
| &bufferRef, | Constructs a host_accessor for accessing a |
| <pre>const property_list &propList = {})</pre> | SYCL buffer immediately on the host. The |
| | optional property_list provides properties |
| | for the constructed SYCL host_accessor |
| | object. |
| <pre>host_accessor(buffer<datat, allocatort="" dimensions,=""></datat,></pre> | Available only when: (dimensions > 0). |
| &bufferRef, | Constructs a host_accessor for accessing a |
| <pre>TagT tag, const property_list &propList = {})</pre> | SYCL buffer immediately on the host. The |
| | tag is used to deduce template arguments |
| | of the host_accessor as described in Sec- |
| | tion 4.7.6.3. The optional property_list |
| | provides properties for the constructed |
| | SYCL host_accessor object. |
| <pre>host_accessor(buffer<datat, allocatort="" dimensions,=""></datat,></pre> | Available only when: (dimensions > 0). |
| &bufferRef, | Constructs a host_accessor for accessing |
| <pre>handler &commandGroupHandlerRef,</pre> | a SYCL buffer inside a host task. The |
| <pre>const property_list &propList = {})</pre> | optional property_list provides properties |
| | for the constructed SYCL host_accessor |
| | object. |
| | Continued on next page |

Table 4.53: Constructors of the host_accessor class template.

| Constructor | Description |
|---|--|
| <pre>host_accessor(buffer<datat, allocatort="" dimensions,=""></datat,></pre> | Available only when: (dimensions > 0). |
| &bufferRef, | Constructs a host_accessor for accessing |
| handler &commandGroupHandlerRef, | a SYCL buffer inside a host task. The |
| <pre>TagT tag, const property_list &propList = {})</pre> | tag is used to deduce template arguments |
| | of the host_accessor as described in Sec- |
| | tion 4.7.6.3. The optional property_list |
| | provides properties for the constructed |
| | SYCL host_accessor object. |
| <pre>host_accessor(buffer<datat, allocatort="" dimensions,=""></datat,></pre> | Available only when: (dimensions > 0). |
| &bufferRef, | Constructs a host_accessor for accessing |
| <pre>range<dimensions> accessRange,</dimensions></pre> | a range of a SYCL buffer immediately on |
| <pre>const property_list &propList = {})</pre> | the host. The optional property_list pro- |
| | vides properties for the constructed SYCL |
| | host_accessor object. |
| <pre>host_accessor(buffer<datat, allocatort="" dimensions,=""></datat,></pre> | Available only when: (dimensions > 0). |
| &bufferRef, | Constructs a host_accessor for accessing |
| <pre>range<dimensions> accessRange,</dimensions></pre> | a range of a SYCL buffer immediately on |
| <pre>TagT tag, const property_list &propList = {})</pre> | the host. The tag is used to deduce tem- |
| | plate arguments of the host_accessor as |
| | described in Section 4.7.6.3. The optional |
| | <pre>property_list provides properties for the</pre> |
| | constructed SYCL host_accessor object. |
| <pre>host_accessor(buffer<datat, allocatort="" dimensions,=""></datat,></pre> | Available only when: (dimensions > 0). |
| &bufferRef, | Constructs a host_accessor for accessing |
| <pre>range<dimensions> accessRange,</dimensions></pre> | a range of a SYCL buffer immediately |
| <pre>id<dimensions> accessOffset,</dimensions></pre> | on the host. The range is specified by |
| <pre>const property_list &propList = {})</pre> | accessRange and accessOffset. The op- |
| | tional property_list provides properties |
| | for the constructed SYCL host_accessor |
| | object. |
| <pre>host_accessor(buffer<datat, allocatort="" dimensions,=""></datat,></pre> | Available only when: (dimensions > 0). |
| &bufferRef, | Constructs a host_accessor for accessing |
| <pre>range<dimensions> accessRange,</dimensions></pre> | a range of a SYCL buffer immediately |
| <pre>id<dimensions> accessOffset,</dimensions></pre> | on the host. The range is specified by |
| <pre>TagT tag, const property_list &propList = {})</pre> | accessRange and accessOffset. The tag |
| | is used to deduce template arguments of |
| | the host_accessor as described in Sec- |
| | tion 4.7.6.3. The optional property_list |
| | provides properties for the constructed |
| | SYCL host_accessor object. |
| <pre>host_accessor(buffer<datat, allocatort="" dimensions,=""></datat,></pre> | Available only when: (dimensions > 0). |
| &bufferRef, | Constructs a host_accessor for accessing |
| handler &commandGroupHandlerRef, | a range of a SYCL buffer inside a host |
| <pre>range<dimensions> accessRange,</dimensions></pre> | task. The optional property_list pro- |
| <pre>const property_list &propList = {})</pre> | vides properties for the constructed SYCL |
| | host_accessor object. |
| | Continued on next page |

Table 4.53: Constructors of the host_accessor class template.

| Constructor | Description |
|---|--|
| <pre>host_accessor(buffer<datat, allocatort="" dimensions,=""></datat,></pre> | Available only when: (dimensions > 0). |
| &bufferRef, | Constructs a host_accessor for accessing |
| handler &commandGroupHandlerRef, | a range of a SYCL buffer inside a host |
| <pre>range<dimensions> accessRange,</dimensions></pre> | task. The tag is used to deduce tem- |
| <pre>TagT tag, const property_list &propList = {})</pre> | plate arguments of the host_accessor as |
| | described in Section 4.7.6.3. The optional |
| | <pre>property_list provides properties for the</pre> |
| | constructed SYCL host_accessor object. |
| <pre>host_accessor(buffer<datat, allocatort="" dimensions,=""></datat,></pre> | Available only when: (dimensions > 0). |
| &bufferRef, | Constructs a host_accessor for accessing a |
| handler &commandGroupHandlerRef, | range of a SYCL buffer inside a host task. |
| <pre>range<dimensions> accessRange,</dimensions></pre> | The range is specified by accessRange and |
| <pre>id<dimensions> accessOffset,</dimensions></pre> | accessOffset. The optional property_list |
| <pre>const property_list &propList = {})</pre> | provides properties for the constructed |
| | SYCL host_accessor object. |
| <pre>host_accessor(buffer<datat, allocatort="" dimensions,=""></datat,></pre> | Available only when: (dimensions > 0). |
| &bufferRef, | Constructs a host_accessor for accessing a |
| <pre>handler &commandGroupHandlerRef,</pre> | range of a SYCL buffer inside a host task. |
| <pre>range<dimensions> accessRange,</dimensions></pre> | The range is specified by accessRange and |
| <pre>id<dimensions> accessOffset,</dimensions></pre> | accessOffset. The tag is used to deduce |
| <pre>TagT tag, const property_list &propList = {})</pre> | template arguments of the host_accessor as |
| | described in Section 4.7.6.3. The optional |
| | <pre>property_list provides properties for the</pre> |
| | constructed SYCL host_accessor object. |
| | End of table |

Table 4.53: Constructors of the host_accessor class template.

| Member function | Description | |
|---|--|--|
| <pre>void swap(host_accessor &other);</pre> | Swaps the contents of the current accessor | |
| | with the contents of other. | |
| <pre>size_type byte_size()const noexcept</pre> | Returns the size in bytes of the re- | |
| | gion of a SYCL buffer that this SYCL | |
| | host_accessor is accessing. | |
| <pre>size_type size()const noexcept</pre> | Returns the number of elements in the re- | |
| | gion of a SYCL buffer that this SYCL | |
| | host_accessor is accessing. | |
| <pre>size_type max_size()const noexcept</pre> | Returns the maximum number of elements | |
| | any accessor of this type would be able to | |
| | access. | |
| <pre>bool empty()const noexcept</pre> | Returns true iff (size()== 0). | |
| Continued on next page | | |

 Table 4.54: Member functions of the host_accessor class template.

| Member function | Description |
|---|---|
| <pre>range<dimensions> get_range()const</dimensions></pre> | Available only when: dimensions > 0. |
| | Returns a range object which represents the |
| | number of elements of dataT per dimension |
| | that this host_accessor may access. |
| | The range object returned must equal to the |
| | range of the buffer this host_accessor is |
| | associated with, unless a range was explic- |
| | itly specified when this host_accessor was |
| | constructed. |
| <pre>id<dimensions> get_offset()const</dimensions></pre> | Available only when: dimensions > 0. |
| | Returns an id object which represents the |
| | starting point in number of elements of |
| | dataT for the range that this host_accessor |
| | may access. |
| | The id object returned must equal to id |
| | {0, 0, 0}, unless an offset was explicitly |
| | specified when this host_accessor was con- |
| | structed. |
| operator reference()const | Available only when: (dimensions == 0). |
| | Returns a reference to the element stored |
| | within the SYCL buffer this SYCL |
| | host_accessor is accessing. |
| <pre>reference operator[](id<dimensions> index)const</dimensions></pre> | Available only when: (dimensions > 0). |
| | Returns a reference to the element stored |
| | within the SYCL buffer this SYCL |
| | host_accessor is accessing at the index |
| | specified by index. |
| <pre>unspecified &operator[](size_t index)const</pre> | Available only when: dimensions > 1 . |
| | Returns an instance of an undefined in- |
| | termediate type representing a SYCL |
| | host_accessor of the same type as this |
| | SYCL host_accessor, with the dimen- |
| | sionality dimensions-1 and containing an |
| | implicit SYCL id with index dimensions |
| | set to index. The intermediate type re- |
| | turned must provide all available subscript |
| | operators which take a size_t parameter |
| | defined by the SYCL host_accessor class |
| | that are appropriate for the type it represents |
| | (including this subscript operator). |
| iterator data()const noexcept | Returns a pointer to the memory this SYCL |
| | host_accessor memory is accessing. |
| iterator begin()const noexcept | Returns an iterator to the first element of the |
| | memory within the access range. |
| iterator end()const noexcept | Returns an iterator that points past the last |
| | element of the memory within the access |
| | range. |
| | Continued on next page |

Table 4.54: Member functions of the host_accessor class template.

| Member function | Description |
|--|--|
| <pre>const_iterator cbegin()const noexcept</pre> | Returns a const iterator to the first element |
| | of the memory within the access range. |
| <pre>const_iterator cend()const noexcept</pre> | Returns a const iterator that points past the last element of the memory within the access |
| | range. |
| | End of table |

Table 4.54: Member functions of the host_accessor class template.

4.7.6.10.2 Host buffer accessor properties

The host_accessor supports the same list of properties as a device buffer accessor listed in 4.7.6.9.2.

4.7.6.11 Local accessor

A local accessor provides access to SYCL runtime allocated shared memory via local memory. A SYCL accessor is considered a local accessor if it has the access target target::local. The memory allocated by a local accessor is non-initialised so it is the user's responsibility to construct and destroy objects explicitly if required. The local memory that is allocated is shared between all work-items of a work-group.

A local accessor does not provide access on the host and the memory can not be copied back to the host.

The data type of a local accessor can be any valid SYCL kernel argument (see Section 4.14.4).

Unless the dimensionality of the local accessor is 0, the size of memory allocated by the SYCL runtime is specified by a range provided on construction. The dimensionality of this range must match the dimensionality of the local accessor.

There are three ways that a SYCL accessor can provide access to the elements of the allocated memory. Firstly by passing a SYCL id instance of the same dimensionality as the SYCL accessor subscript operator. Secondly by passing a single size_t value to multiple consecutive subscript operators (one for each dimension of the SYCL accessor, for example acc[z][y][x]). Finally, in the case of the SYCL accessor having 0 dimensions, by triggering the implicit conversion operator. Whenever a multi-dimensional index is passed to a SYCL accessor, the linear index is calculated based on the index {id0, id1, id2} provided and the range of the SYCL accessor {r0, r1, r2} according to row-major ordering as follows:

$$id2 + (id1 \cdot r2) + (id0 \cdot r2 \cdot r1) \tag{4.5}$$

A local accessor can optionally provide atomic access to allocated memory, using the access_mode ::atomic, in which case all operators which return an element of the allocated memory return an instance of the SYCL atomic class. This is deprecated in SYCL 2020.

Local accessors are not valid in the single_task or basic parallel_for SYCL kernel function invocations, due the fact that local work-groups are implicitly created, and the implementation is free to choose any size.

A local accessor meets the C++ requirement of ContiguousContainer. The iterator for this container is multi_ptr <dataT, access::address_space::local_space, access::decorated::no>. For multidimensional accessors the iterator linearizes the data according to 4.3.

The full list of capabilities that local accessors can support is described in 4.55.

| Access target | Accessor type | Access modes | Data types | Dimensionalities | Placeholder |
|---------------|------------------|----------------------|---|--|-------------|
| local | device | read_write atomic | All available data types supported in a SYCL kernel function. | Between 0 and 3 (in- clusive). | No |

Table 4.55: Description of all the local accessor capabilities.

4.7.6.11.1 Local accessor interface

A synopsis of the SYCL accessor class template local specialization is provided below. The member types for this accessor specialization are listed in Tables 4.56. The constructors for this accessor specialization are listed in Tables 4.57. The member functions for this accessor specialization are listed in Tables 4.58. The additional common special member functions and common member functions are listed in 4.5.3 in Tables 4.1 and 4.2, respectively. For valid implicit conversions between accessor types please refer to 4.7.6.8. Additionally, local accessors of the same type must be equality comparable.

```
namespace sycl {
 1
 2
    template <typename dataT,</pre>
 3
              int dimensions,
 4
              access::mode accessmode,
 5
              access::target accessTarget,
              access::placeholder isPlaceholder>
 6
 7
    class accessor {
 8
     public:
 9
      template <typename value_type, access::decorated IsDecorated>
10
      using accessor_ptr = multi_ptr<value_type, access::address_space::local_space, IsDecorated>;
11
      using value_type = dataT;
12
      using reference = dataT &;
13
      using const_reference = const dataT &;
14
      using iterator = accessor_ptr<dataT, access::decorated::no>;
15
      using const_iterator = accessor_ptr<const dataT, access::decorated::no>;
16
      using reverse_iterator = std::reverse_iterator<iterator>;
17
      using const_reverse_iterator = std::reverse_iterator<const_iterator>;
18
      using difference_type =
19
          typename std::iterator_traits<iterator>::difference_type;
20
      using size_type = size_t;
21
22
      accessor();
23
24
      /* Available only when: dimensions == 0 */
25
      accessor(handler &commandGroupHandlerRef,
26
               const property_list &propList = {});
27
28
      /* Available only when: dimensions > 0 */
29
      accessor(range<dimensions> allocationSize, handler &commandGroupHandlerRef,
30
               const property_list &propList = {});
31
32
      /* -- common interface members -- */
33
34
      void swap(accessor &other);
35
36
      size_type byte_size() const noexcept;
37
```

38 size_type size() const noexcept; 39 40 size_type max_size() const noexcept; 41 42 size_type get_count() const noexcept; 43 44 bool empty() const noexcept; 45 46 range<dimensions> get_range() const; 47 48 /* Available only when: (dimensions == 0) */ 49 operator reference() const; 50 /* Available only when: (dimensions > 0) */ 51 52 reference operator[](id<dimensions> index) const; 53 54 /* Deprecated in SYCL 2020 Available only when: accessMode == access::mode::atomic && dimensions == 0 */ 55 56 operator cl::sycl::atomic<dataT,access::address_space::local_space> () const; 57 58 /* Deprecated in SYCL 2020 59 Available only when: accessMode == access::mode::atomic && dimensions > 0 */ 60 cl::sycl::atomic<dataT, access::address_space::local_space> operator[](61 id<dimensions> index) const; 62 /* Available only when: dimensions > 1 */ 63 64 __unspecified__ &operator[](size_t index) const; 65 66 std::add_pointer_t<value_type> get_pointer() const noexcept; 67 68 template <access::decorated IsDecorated> 69 accessor_ptr<value_type, IsDecorated> get_multi_ptr() const noexcept; 70 71 iterator data() const noexcept; 72 73 iterator begin() const noexcept; 74 75 iterator end() const noexcept; 76 77 const_iterator cbegin() const noexcept; 78 79 const_iterator cend() const noexcept; 80 81 reverse_iterator rbegin() const noexcept; 82 83 reverse_iterator rend() const noexcept; 84 85 const_reverse_iterator crbegin() const noexcept; 86 87 const_reverse_iterator crend() const noexcept; 88 }; } // namespace sycl 89

Listing 4.5: Accessor class for locals.

| Member types | Description |
|--|---|
| <pre>template <access::decorated isdecorated=""></access::decorated></pre> | <pre>multi_ptr<value_type, access::<="" pre=""></value_type,></pre> |
| accessor_ptr | <pre>address_space::local_space,</pre> |
| | IsDecorated>. |
| value_type | dataT |
| reference | dataT& |
| const_reference | const dataT& |
| iterator | <pre>raw_local_ptr<value_type></value_type></pre> |
| const_iterator | <pre>raw_local_ptr<const value_type=""></const></pre> |
| reverse_iterator | Iterator adaptor that reverses the direction |
| | of iterator. |
| const_reverse_iterator | Iterator adaptor that reverses the direction |
| | of const_iterator. |
| difference_type | <pre>typename std::iterator_traits<</pre> |
| | <pre>iterator>::difference_type</pre> |
| size_type | size_t |
| | End of table |

Table 4.56: Member types of the accessor class template local specialization .

| Constructor | Description |
|--|--|
| accessor() | Constructs an empty accessor. Fulfills the |
| | following post-conditions: |
| | <pre>• (empty()== true)</pre> |
| | All size queries return 0. |
| | • The only iterator that can be obtained |
| | is nullptr. |
| | • Trying to access the underlying mem- |
| | ory is undefined behavior. |
| | |
| <pre>accessor(handler &commandGroupHandlerRef,</pre> | Available only when: dimensions == 0. |
| <pre>const property_list &propList = {})</pre> | Constructs a SYCL accessor instance for |
| | accessing runtime allocated local memory of |
| | a single element (of type dataT) within a |
| | SYCL kernel function on the SYCL queue |
| | associated with commandGroupHandlerRef. |
| | The allocation is per work-group. The |
| | optional property_list provides properties |
| | for the constructed SYCL accessor object. |
| | Continued on next page |

 Table 4.57: Constructors of the accessor class template local specialization.

| Constructor | Description |
|--|--|
| <pre>accessor(range<dimensions> allocationSize,</dimensions></pre> | Available only when: dimensions > 0 . |
| handler &commandGroupHandlerRef, | Constructs a SYCL accessor instance for |
| <pre>const property_list &propList = {})</pre> | accessing runtime allocated local memory of |
| | size specified by allocationSize within a |
| | SYCL kernel function on the SYCL queue |
| | associated with commandGroupHandlerRef. |
| | allocationSize defines the number of ele- |
| | ments of type dataT to be allocated. The al- |
| | location is per work-group, and if multiple |
| | work-groups execute simultaneously in an |
| | implementation, each work-group will re- |
| | ceive its own functionally independent allo- |
| | cation of size allocationSize elements of |
| | type dataT. The optional property_list |
| | provides properties for the constructed |
| | SYCL accessor object. |
| | End of table |

Table 4.57: Constructors of the accessor class template local specialization.

| Member function | Description |
|--|--|
| <pre>void swap(accessor &other);</pre> | Swaps the contents of the current accessor |
| | with the contents of other. |
| <pre>size_type byte_size()const noexcept</pre> | Returns the size in bytes of the local mem- |
| | ory allocation, per work-group, that this |
| | SYCL accessor is accessing. |
| <pre>size_type size()const noexcept</pre> | Returns the number of dataT elements |
| | in the local memory allocation, per work- |
| | group, that this SYCL accessor is access- |
| | ing. |
| <pre>size_type max_size()const noexcept</pre> | Returns the maximum number of elements |
| | any accessor of this type would be able to |
| | access. |
| <pre>size_type get_count()const noexcept</pre> | Returns the same as size(). |
| <pre>bool empty()const noexcept</pre> | Returns true iff (size()== 0). |
| <pre>range<dimensions> get_range()const</dimensions></pre> | Available only when: dimensions > 0 . |
| | Returns a range object which represents the |
| | number of elements of dataT per dimension |
| | that this accessor may access, per work- |
| | group. |
| operator reference()const | Available only when: (dimensions == 0). |
| | Returns a reference to the single element |
| | stored within the work-group's local mem- |
| | ory allocation that this accessor is access- |
| | ing. |
| | Continued on next page |

Table 4.58: Member functions of the accessor class template local specialization.

| Member function | Description |
|---|---|
| <pre>reference operator[](id<dimensions> index)const</dimensions></pre> | Available only when: (dimensions > 0). |
| | Returns a reference to the element stored |
| | within the work-group's local memory allo- |
| | cation that this SYCL accessor is accessing, |
| | at the index specified by index. |
| <pre>operator cl::sycl::atomic<datat,< pre=""></datat,<></pre> | Deprecated in SYCL 2020. |
| <pre>access::address_space::local_space> &()const</pre> | Available only when: accessMode == |
| | <pre>access_mode::atomic && dimensions ==</pre> |
| | 0). |
| | Returns a reference to an instance of cl |
| | ::sycl::atomic of type dataT providing |
| | atomic access to the element stored within |
| | the work-group's local memory allocation |
| | that this SYCL accessor is accessing. |
| <pre>cl::sycl::atomic<datat, access::address_space::<="" pre=""></datat,></pre> | Deprecated in SYCL 2020. |
| local_space> & | Available only when: accessMode == |
| operator[](1d <dimensions> index)const</dimensions> | access_mode::atomic && dimensions > |
| | 0). Deturns a reference to an instance of cl |
| | result is a reference to an instance of er |
| | atomic access to the element stored within |
| | the work-group's local memory allocation |
| | that this SYCL accessor is accessing, at the |
| | index specified by index. |
| unspecified & operator[](size t index)const | Available only when: dimensions > 1 . |
| | Returns an instance of an undefined inter- |
| | mediate type representing a SYCL accessor |
| | of the same type as this SYCL accessor, |
| | with the dimensionality dimensions-1 and |
| | containing an implicit SYCL id with index |
| | dimensions set to index. The intermedi- |
| | ate type returned must provide all available |
| | subscript operators which take a size_t pa- |
| | rameter defined by the SYCL accessor class |
| | that are appropriate for the type it represents |
| | (including this subscript operator). |
| <pre>std::add_pointer_t<value_type> get_pointer()const</value_type></pre> | Returns a pointer to the work-group's local |
| noexcept | memory allocation that this SYCL accessor |
| | is accessing. |
| template <access::decorated isdecorated=""></access::decorated> | Returns a multi_ptr to the work-group's |
| accessor_ptr <isdecorated> get_multi_ptr()const</isdecorated> | local memory allocation that this SYCL |
| noexcept | accessor is accessing. |
| iterator data()const noexcept | memory allocation that this SVCI |
| | is accessing |
| iterator hegin() const nearcont | Returns an iterator to the first element of |
| | allocated local memory |
| | Continued on payt page |

Table 4.58: Member functions of the accessor class template local specialization.

| Member function | Description |
|---|---|
| iterator end()const noexcept | Returns an iterator that points past the last |
| | element of allocated local memory. |
| <pre>const_iterator cbegin()const noexcept</pre> | Returns a const iterator to the first element |
| | of allocated local memory. |
| <pre>const_iterator cend()const noexcept</pre> | Returns a const iterator that points past the |
| | last element of allocated local memory. |
| <pre>reverse_iterator rbegin()const noexcept</pre> | Returns an iterator adaptor to the last ele- |
| | ment of the memory of allocated local mem- |
| | ory. |
| reverse_iterator rend()const noexcept | Returns an iterator adaptor that points be- |
| | fore the first element of the memory of allo- |
| | cated local memory. |
| <pre>const_reverse_iterator crbegin()const noexcept</pre> | Returns a const iterator adaptor to the last |
| | element of the memory of allocated local |
| | memory. |
| <pre>const_reverse_iterator crend()const noexcept</pre> | Returns a const iterator adaptor that points |
| | before the first element of the memory of al- |
| | located local memory. |
| | End of table |

Table 4.58: Member functions of the accessor class template local specialization.

4.7.6.11.2 Local accessor properties

The property_list constructor parameters are present for extensibility.

4.7.6.12 Image accessor

An image accessor provides access to either an instance of a SYCL unsampled_image or sampled_image. A SYCL accessor is considered an image accessor if it has the access target target::unsampled_image, target::sampled_image, target::host_unsampled_image or target::host_sampled_image.

An image accessor can provide access to memory managed by a SYCL unsampled_image or sampled_image class, using the access target target::unsampled_image or target::sampled_image.

Alternatively an image accessor can provide access to memory managed by a SYCL unsampled_image or sampled_image immediately on the host, using the access target target::host_unsampled_image or target:: host_sampled_image, respectively. If the SYCL image this SYCL accessor is accessing was constructed with the property property::image::use_host_ptr the address of the memory accessed on the host must be the address the SYCL image was constructed with, otherwise the SYCL runtime is free to allocate temporary memory to provide access on the host.

The data type of an image accessor must be either int4, uint4, float4 or half4.

The dimensionality of an image accessor must match that of the SYCL image which it is providing access to, with the exception of when the access target is target::image_array, in which case the dimensionality of the SYCL accessor must be 1 less.

An image accessor with the access target target::image or target::host_image can provide access to the elements of a SYCL image by passing a SYCL int4 or float4 instance to the read or write member functions.

An image accessor with the access target target::image_array can provide access to a slice of an image array by passing a size_t value to the subscript operator. This returns an instance of __image_array_slice__, an unspecified type providing the interface of accessor<dataT, dimensions, mode, target::image> which will provide access to a slice of the image array specified by index. The __image_array_slice__ returned can then provide access via the read or write member functions as described above. For example acc[arrayIndex].read (coords).

The full list of capabilities that image accessors can support is described in 4.59.

| Access target | Accessor type | Access modes | Data types | Dimensionalities | Placeholder |
|-------------------------|------------------|--------------------------------|----------------------------------|-----------------------------------|-------------|
| unsampled_image | device | read write discard_write | int4 uint4 float4 half4 | Between 1 and 3 (in- clusive). | No |
| sampled_image | device | read | int4 uint4 float4 half4 | Between 1 and 3 (in- clusive). | No |
| host unsampled_image | host | read write discard_write | int4 uint4 float4 half4 | Between 1 and 3 (in- clusive). | No |
| host_sampled image | host | read | int4 uint4 float4 half4 | Between 1 and 3 (in- clusive). | No |

Table 4.59: Description of all the image accessor capabilities.

4.7.6.12.1 Image accessor interface

A synopsis of the SYCL accessor class template image specialization is provided below. The constructors and member functions of the SYCL accessor class template image specialization are listed in Tables 4.60 and 4.61 respectively. The additional common special member functions and common member functions are listed in 4.5.3 in Tables 4.1 and 4.2, respectively. For valid implicit conversions between accessor types please refer to 4.7.6.8.

```
1
    namespace sycl {
 2
 3
    template <typename dataT,</pre>
 4
              int dimensions,
 5
              access::mode accessMode,
 6
              access::target accessTarget>
 7
    class accessor {
 8
     public:
 9
      using value_type = dataT;
10
      using reference = dataT &;
11
      using const_reference = const dataT &;
12
13
      /* Available only when: accessTarget == access::target::host_unsampled_image */
14
      template <typename AllocatorT>
15
      accessor(unsampled_image<dimensions, AllocatorT> & imageRef);
16
```

```
17
      /* Available only when: accessTarget == access::target::host_sampled_image */
18
      template <typename AllocatorT>
19
      accessor(sampled_image<dimensions, AllocatorT> & imageRef);
20
21
      /* Available only when: accessTarget == access::target::unsampled_image */
22
      template <typename AllocatorT>
23
      accessor(unsampled_image<dimensions, AllocatorT> & imageRef,
24
       handler &commandGroupHandlerRef);
25
26
      /* Available only when: accessTarget == access::target::sampled_image */
27
      template <typename AllocatorT>
      accessor(sampled_image<dimensions, AllocatorT> & imageRef,
28
29
       handler &commandGroupHandlerRef);
30
31
      /* -- common interface members -- */
32
33
     /* -- property interface members -- */
34
35
     size_t get_count() const;
36
37
      /* Available only when: (accessTarget == access::target::unsampled_image &&
      accessMode == access::mode::read) || (accessTarget ==
38
39
      access::target::host_unsampled_image && accessMode == access::mode::read)
40
      if dimensions == 1, coordT = int
41
      if dimensions == 2, coordT = int2
42
      if dimensions == 4, coordT = int4 */
43
      template <typename coordT>
44
      dataT read(const coordT &coords) const noexcept;
45
46
      /* Available only when: (accessTarget == access::target::sampled_image &&
47
      accessMode == access::mode::read) || (accessTarget ==
48
      access::target::host_sampled_image && accessMode == access::mode::read)
49
      if dimensions == 1, coordT = float
50
      if dimensions == 2, coordT = float2
51
      if dimensions == 3, coordT = float4 */
      template <typename coordT>
52
53
      dataT read(const coordT &coords) const noexcept;
54
55
      /* Available only when: (accessTarget == access::target::unsampled_image &&
56
      (accessMode == access::mode::write || accessMode == access::mode::discard_write)) ||
      (accessTarget == access::target::host_unsampled_image && (accessMode == access::mode::write ||
57
58
      accessMode == access::mode::discard_write))
59
     if dimensions == 1, coordT = int
60
     if dimensions == 2, coordT = int2
61
     if dimensions == 3, coordT = int4 */
62
     template <typename coordT>
63
     void write(const coordT &coords, const dataT &color) const;
64 };
65 } // namespace sycl
                                 Listing 4.6: Accessor interface for images.
```

| Constructor | Description |
|--|--|
| <pre>template <typename allocatort=""></typename></pre> | Available only when: accessTarget == |
| <pre>accessor(unsampled_image<dimensions, allocatort<="" pre=""></dimensions,></pre> | <pre>target::host_unsampled_image.</pre> |
| >, | Constructs a SYCL accessor instance for |
| <pre>&imageRef, const property_list &propList = {})</pre> | accessing a SYCL unsampled_image im- |
| | mediately on the host. The optional |
| | <pre>property_list provides properties for the</pre> |
| | constructed SYCL accessor object. |
| <pre>template <typename allocatort=""></typename></pre> | Available only when: accessTarget == |
| <pre>accessor(sampled_image<dimensions, allocatort="">,</dimensions,></pre> | <pre>target::host_sampled_image.</pre> |
| &imageRef, const property_list &propList = {}) | Constructs a SYCL accessor instance for |
| | accessing a SYCL sampled_image im- |
| | mediately on the host. The optional |
| | <pre>property_list provides properties for the</pre> |
| | constructed SYCL accessor object. |
| <pre>template <typename allocatort=""></typename></pre> | Available only when: accessTarget == |
| <pre>accessor(unsampled_image<dimensions, allocatort<="" pre=""></dimensions,></pre> | <pre>target::unsampled_image.</pre> |
| >, | Constructs a SYCL accessor instance for |
| <pre>&imageRef, handler &commandGroupHandlerRef,</pre> | accessing a SYCL unsampled_image within |
| <pre>const property_list &propList = {})</pre> | a SYCL kernel function on the SYCL queue |
| | associated with commandGroupHandlerRef. |
| | The optional property_list provides prop- |
| | erties for the constructed SYCL accessor |
| | object. |
| <pre>template <typename allocatort=""></typename></pre> | Available only when: accessTarget == |
| accessor(sampled_image <dimensions, allocatort="">,</dimensions,> | target::sampled_image. |
| &imageRef, handler &commandGroupHandlerRef, | Constructs a SYCL accessor instance for |
| <pre>const property_list &propList = {})</pre> | accessing a SYCL sampled_image within a |
| | SYCL kernel function on the SYCL queue |
| | The optional property list provides pro- |
| | article for the constructed SVCL |
| | object |
| | Find of table |
| | End of table |

Table 4.60: Constructors of the accessor class template image specialization.

| Member function | Description |
|------------------------------------|---------------------------------------|
| <pre>size_t get_size()const</pre> | Returns the size in bytes of the SYCL |
| | unsampled_image or sampled_image this |
| | SYCL accessor is accessing. |
| <pre>size_t get_count()const</pre> | Returns the number of elements of the |
| | SYCL unsampled_image or sampled_image |
| | this SYCL accessor is accessing. |
| | Continued on next page |

Table 4.61: Member functions of the **accessor** class template image specialization.

| Member function | Description |
|---|---|
| <pre>template <typename coordt=""></typename></pre> | Available only when: (accessTarget |
| <pre>dataT read(const coordT &coords)const</pre> | == target::unsampled_image && |
| | <pre>accessMode == access_mode::read</pre> |
| |) (accessTarget == target:: |
| | host_unsampled_image && accessMode |
| | <pre>== access_mode::read).</pre> |
| | Reads and returns an element of the |
| | unsampled_image at the coordinates speci- |
| | fied by coords. Permitted types for coordT |
| | are int when dimensions == 1, int2 |
| | when dimensions == 2 and int4 when |
| | dimensions == 3. |
| <pre>template <typename coordt=""></typename></pre> | Available only when: (accessTarget == |
| <pre>dataT read(const coordT &coords)const</pre> | <pre>target::sampled_image && accessMode</pre> |
| | == access_mode::read) (accessTarget |
| | <pre>== target::host_sampled_image &&</pre> |
| | <pre>accessMode == access_mode::read).</pre> |
| | Reads and returns a sampled element of the |
| | <pre>sampled_image at the coordinates specified</pre> |
| | by coords. Permitted types for coordT |
| | are float when dimensions == 1, float2 |
| | when dimensions $=$ 2 and float4 when |
| | dimensions == 3. |
| <pre>template <typename coordt=""></typename></pre> | Available only when: (accessTarget |
| <pre>void write(const coordT &coords, const dataT &color)</pre> | == target::unsampled_image && (|
| const | <pre>accessMode == access_mode::write</pre> |
| | <pre> accessMode == access_mode::</pre> |
| | discard_write)) (accessTarget == |
| | <pre>target::host_unsampled_image &&</pre> |
| | <pre>(accessMode == access_mode::write</pre> |
| | <pre> accessMode == access_mode::</pre> |
| | discard_write)). |
| | Writes the value specified by color to the |
| | element of the image at the coordinates |
| | specified by coords. Permitted type for |
| | coordT are int when dimensions == 1, |
| | int2 when dimensions $=$ 2 and int4 |
| | when dimensions == 3. |
| image_array_slice | Available only when: accessTarget == |
| operator[](size_t index)const | <pre>target::image_array && dimensions < 3</pre> |
| | Determine |
| | Returns an instance of |
| | an unspeci- |
| | neu type which provides the interface of |
| | accessor< data1, dimensions, mode, |
| | ta a alian of the impact arrest area: 6, 1, 1 |
| | index |
| | inaex. |
| | End of table |

Table 4.61: Member functions of the accessor class template image specialization.

4.7.6.12.2 Image accessor properties

The property_list constructor parameters are present for extensibility.

4.7.7 Address space classes

In SYCL, there are five different address spaces: global, local, constant, private and generic. In a SYCL generic implementation, types are not affected by the address spaces. However, there are situations where users need to explicitly carry address spaces in the type. For example:

- For performance tuning and genericity. Even if the platform supports the representation of the generic address space, this may come at some performance sacrifice. In order to help the target compiler, it can be useful to track specifically which address space a pointer is addressing.
- When linking SYCL kernels with SYCL backend-specific functions. In this case, it might be necessary to specify the address space for any pointer parameters.

Direct declaration of pointers with address spaces is discouraged as the definition is implementation defined. Users must rely on the multi_ptr class to handle address space boundaries and interoperability.

4.7.7.1 Multi-pointer class

The multi-pointer class is the common interface for the explicit pointer classes, defined in 4.7.7.2.

There are situations where a user may want to make their type address space dependent. This allows performing generic programming that depends on the address space associated with their data. An example might be wrapping a pointer inside a class, where a user may need to template the class according to the address space of the pointer the class is initialized with. In this case, the multi_ptr class enables users to do this in a portable and stable way.

The multi_ptr class exposes 3 flavors of the same interface. If the value of access::decorated is access:: decorated::no, the interface exposes pointers and references type that are not decorated by an address space. If the value of access::decorated is access::decorated::yes, the interface exposes pointers and references type that are decorated by an address space. The decorated::yes, the interface exposes pointers and references type extensions. The decorated type may be distinct from the non-decorated one. For interoperability with the SYCL backend, users should rely on types exposed by the decorated version. If the value of access::decorated is access::decorated is access::decorated is access::decorated version.

The template traits remove_decoration and type alias remove_decoration_t retrieve the non-decorated pointer or reference from a decorated one. Using this template trait with a non-decorated type is safe and returns the same type.

It is possible to use the void type for the multi_ptr class, but in that case some functionality is disabled. multi_ptr<void> does not provide the reference or const_reference types, the access operators (operator *(), operator->()), the arithmetic operators or prefetch member function. Conversions from multi_ptr to multi_ptr<void> of the same address space are allowed, and will occur implicitly. Conversions from multi_ptr <void> to any other multi_ptr type of the same address space are allowed, but must be explicit. The same rules apply to multi_ptr<const void>.

An overview of the interface provided for the multi_ptr class follows.

```
1 namespace sycl {
2 namespace access {
```

```
enum class address_space : int {
3
 4
     global_space,
5
     local_space,
6
     constant_space,
7
     private_space,
8
     generic_space,
9 };
10
11 enum class decorated : int {
12
     no.
13
     yes,
14
     legacy,
15 };
16
17 } // namespace access
18
19 template<typename T> struct remove_decoration {
20
    using type = /* ... */;
21 };
22
23 template<typename T>
24 using remove_decoration_t = remove_decoration::type;
25
26 template <typename ElementType, access::address_space Space, access::decorated DecorateAddress>
27
   class multi_ptr {
28
    public:
29
     static constexpr bool is_decorated = DecorateAddress == access::decorated::yes;
30
     static constexpr access::address_space address_space = Space;
31
32
     using value_type = ElementType;
33
     using pointer = std::conditional<is_decorated, __unspecified__ *,</pre>
34
                                       std::add_pointer_t<value_type>>;
35
     using reference = std::conditional<is_decorated, __unspecified__ &,</pre>
36
                                         std::add_lvalue_reference_t<value_type>>;
37
     using iterator_category = std::random_access_iterator_tag;
38
     using difference_type = std::ptrdiff_t;
39
40
      static_assert(std::is_same_v<remove_decoration_t<pointer>, std::add_pointer_t<value_type>>>);
41
      static_assert(std::is_same_v<remove_decoration_t<reference>, std::add_lvalue_reference_t<</pre>
          value_type>>);
42
      // Legacy has a different interface.
43
      static_assert(DecorateAddress != access::decorated::legacy);
44
45
     // Constructors
46
     multi_ptr();
47
     multi_ptr(const multi_ptr&);
48
     multi_ptr(multi_ptr&&);
49
      explicit multi_ptr(multi_ptr<ElementType, Space, yes>::pointer);
50
     multi_ptr(std::nullptr_t);
51
52
     // Only if Space == global_space or generic_space
53
      template <int dimensions, access::mode Mode, access::placeholder isPlaceholder>
54
     multi_ptr(accessor<value_type, dimensions, Mode, access::target::global_buffer, isPlaceholder>);
55
56
      // Only if Space == local_space or generic_space
```

```
57
       template <int dimensions, access::mode Mode, access::placeholder isPlaceholder>
58
       multi_ptr(accessor<value_type, dimensions, Mode, access::target::local, isPlaceholder>);
59
60
       // Only if Space == constant_space
61
       template <int dimensions, access::mode Mode, access::placeholder isPlaceholder>
62
       multi_ptr(accessor<value_type, dimensions, Mode, access::target::constant_buffer, isPlaceholder</pre>
           >);
63
       // Assignment and access operators
64
65
      multi_ptr &operator=(const multi_ptr&);
66
      multi_ptr &operator=(multi_ptr&&);
67
      multi_ptr &operator=(std::nullptr_t);
68
69
       // Only if Space == address_space::generic_space
70
       // and ASP != access::address_space::constant_space
71
       template<access::address_space ASP, access::decorated IsDecorated>
72
      multi_ptr &operator=(const multi_ptr<value_type, ASP, IsDecorated>&);
73
      // Only if Space == address_space::generic_space
74
       // and ASP != access::address_space::constant_space
75
       template<access::address_space ASP, access::decorated IsDecorated>
      multi_ptr &operator=(multi_ptr<value_type, ASP, IsDecorated>&&);
76
77
78
      reference operator*() const;
79
      pointer operator->() const;
80
81
       pointer get() const;
82
       std::add_pointer_t<value_type> get_raw() const;
83
       __unspecified__ * get_decorated() const;
84
85
       // Conversion to the underlying pointer type
86
       // Deprecated, get() should be used instead.
87
      operator pointer() const;
88
89
      // Only if Space == address_space::generic_space
90
      // Cast to private_ptr
91
      explicit operator multi_ptr<value_type, access::address_space::private_space,</pre>
92
                                   DecorateAddress>();
93
       // Only if Space == address_space::generic_space
94
       // Cast to private_ptr
95
       explicit
96
       operator multi_ptr<const value_type, access::address_space::private_space,</pre>
97
                          DecorateAddress>() const;
98
       // Only if Space == address_space::generic_space
99
       // Cast to global_ptr
100
       explicit operator multi_ptr<value_type, access::address_space::global_space,
101
                                   DecorateAddress>();
102
      // Only if Space == address_space::generic_space
103
       // Cast to global_ptr
104
      explicit
105
       operator multi_ptr<const value_type, access::address_space::global_space,</pre>
106
                          DecorateAddress>() const;
107
       // Only if Space == address_space::generic_space
108
       // Cast to local_ptr
109
       explicit operator multi_ptr<value_type, access::address_space::local_space,</pre>
110
                                   DecorateAddress>();
```
111 // Only if Space == address_space::generic_space // Cast to local_ptr 112 113 explicit 114 operator multi_ptr<const value_type, access::address_space::local_space,</pre> 115 DecorateAddress>() const; 116 117 // Implicit conversion to a multi_ptr<void>. // Only available when value_type is not const-gualified. 118 119 template<access::decorated DecorateAddress> 120 operator multi_ptr<void, Space, DecorateAddress>() const; 121 122 // Implicit conversion to a multi_ptr<const void>. 123 // Only available when value_type is const-qualified. 124 template<access::decorated DecorateAddress> 125 operator multi_ptr<const void, Space, DecorateAddress>() const; 126 127 // Implicit conversion to multi_ptr<const value_type, Space>. 128 template<access::decorated DecorateAddress> 129 operator multi_ptr<const value_type, Space, DecorateAddress>() const; 130 131 // Implicit conversion to the non-decorated version of multi_ptr. 132 // Only available when is_decorated is true. 133 operator multi_ptr<value_type, Space, access::decorated::no>() const; 134 135 // Implicit conversion to the decorated version of multi_ptr. 136 // Only available when is_decorated is false. 137 operator multi_ptr<value_type, Space, access::decorated::yes>() const; 138 139 void prefetch(size_t numElements) const; 140 141 // Arithmetic operators 142 friend multi_ptr& operator++(multi_ptr& mp) { /* ... */ } 143 friend multi_ptr operator++(multi_ptr& mp, int) { /* ... */ } 144 friend multi_ptr& operator--(multi_ptr& mp) { /* ... */ } friend multi_ptr operator--(multi_ptr& mp, int) { /* ... */ } 145 146 friend multi_ptr& operator+=(multi_ptr& lhs, difference_type r) { /* ... */ } 147 friend multi_ptr& operator-=(multi_ptr& lhs, difference_type r) { /* ... */ } 148 friend multi_ptr operator+(const multi_ptr& lhs, difference_type r) { /* ... */ } friend multi_ptr operator-(const multi_ptr& lhs, difference_type r) { /* ... */ } 149 150 friend reference operator*(const multi_ptr& lhs) { /* ... */ } 151 friend bool operator==(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ } 152 153 friend bool operator!=(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ } 154 friend bool operator<(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }</pre> 155 friend bool operator>(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ } friend bool operator<=(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ } 156 157 friend bool operator>=(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ } 158 159 friend bool operator==(const multi_ptr& lhs, std::nullptr_t) { /* ... */ } friend bool operator!=(const multi_ptr& lhs, std::nullptr_t) { /* ... */ } 160 friend bool operator<(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }</pre> 161 friend bool operator>(const multi_ptr& lhs, std::nullptr_t) { /* ... */ } 162 friend bool operator<=(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }</pre> 163 164 friend bool operator>=(const multi_ptr& lhs, std::nullptr_t) { /* ... */ } 165

```
166
       friend bool operator==(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }
167
       friend bool operator!=(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }
168
       friend bool operator<(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }</pre>
169
       friend bool operator>(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }
       friend bool operator<=(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }</pre>
170
171
       friend bool operator>=(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }
172
173 };
174
175 // Specialization of multi_ptr for void and const void
176 // VoidType can be either void or const void
     template <access::address_space Space, access::decorated DecorateAddress>
177
178 class multi_ptr<VoidType, Space, DecorateAddress> {
179
      public:
180
       static constexpr bool is_decorated = DecorateAddress == access::decorated::yes;
181
       static constexpr access::address_space address_space = Space;
182
183
       using value_type = VoidType;
184
       using pointer = std::conditional<is_decorated, __unspecified__ value_type *,</pre>
185
                                         std::add_pointer_t<value_type>>;
186
       using difference_type = std::ptrdiff_t;
187
188
       static_assert(std::is_same_v<remove_decoration_t<pointer>, std::add_pointer_t<value_type>>);
       // Legacy has a different interface.
189
190
       static_assert(DecorateAddress != access::decorated::legacy);
191
192
       // Constructors
193
       multi_ptr();
194
       multi_ptr(const multi_ptr&);
195
       multi_ptr(multi_ptr&&);
196
       explicit multi_ptr(multi_ptr<VoidType, Space, yes>::pointer);
197
       multi_ptr(std::nullptr_t);
198
199
       // Only if Space == global_space
       template <typename ElementType, int dimensions, access::mode Mode,</pre>
200
                 access::placeholder isPlaceholder>
201
202
       multi_ptr(accessor<ElementType, dimensions, Mode,</pre>
203
                          access::target::global_buffer, isPlaceholder>);
204
205
       // Only if Space == local_space
206
       template <typename ElementType, int dimensions, access::mode Mode,</pre>
207
                 access::placeholder isPlaceholder>
208
       multi_ptr(accessor<ElementType, dimensions, Mode, access::target::local,</pre>
209
                          isPlaceholder>);
210
211
       // Only if Space == constant_space
212
       template <typename ElementType, int dimensions, access::mode Mode,</pre>
                 access::placeholder isPlaceholder>
213
214
       multi_ptr(accessor<ElementType, dimensions, Mode,</pre>
215
                          access::target::constant_buffer, isPlaceholder>);
216
217
       // Assignment operators
218
       multi_ptr &operator=(const multi_ptr&);
219
       multi_ptr &operator=(multi_ptr&&);
220
       multi_ptr &operator=(std::nullptr_t);
```

221 222 pointer get() const; 223 224 // Conversion to the underlying pointer type 225 explicit operator pointer() const; 226 227 // Explicit conversion to a multi_ptr<ElementType> 228 // If VoidType is const, ElementType must be as well 229 template <typename ElementType> 230 explicit operator multi_ptr<ElementType, Space, DecorateAddress>() const; 231 232 // Implicit conversion to the non-decorated version of multi_ptr. 233 // Only available when is_decorated is true. 234 operator multi_ptr<value_type, Space, access::decorated::no>() const; 235 236 // Implicit conversion to the decorated version of multi_ptr. 237 // Only available when is_decorated is false. 238 operator multi_ptr<value_type, Space, access::decorated::yes>() const; 239 240 // Implicit conversion to multi_ptr<const void, Space> 241 operator multi_ptr<const void, Space, DecorateAddress>() const; 242 243 friend bool operator==(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ } friend bool operator!=(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ } 244 friend bool operator<(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }</pre> 245 246 friend bool operator>(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ } 247 friend bool operator<=(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ } 248 friend bool operator>=(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ } 249 250 friend bool operator==(const multi_ptr& lhs, std::nullptr_t) { /* ... */ } 251 friend bool operator!=(const multi_ptr& lhs, std::nullptr_t) { /* ... */ } 252 friend bool operator<(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }</pre> friend bool operator>(const multi_ptr& lhs, std::nullptr_t) { /* ... */ } 253 254 friend bool operator<=(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }</pre> 255 friend bool operator>=(const multi_ptr& lhs, std::nullptr_t) { /* ... */ } 256 257 friend bool operator==(std::nullptr_t, const multi_ptr& rhs) { /* ... */ } 258 friend bool operator!=(std::nullptr_t, const multi_ptr& rhs) { /* ... */ } friend bool operator<(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }</pre> 259 260 friend bool operator>(std::nullptr_t, const multi_ptr& rhs) { /* ... */ } 261 friend bool operator<=(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }</pre> 262 friend bool operator>=(std::nullptr_t, const multi_ptr& rhs) { /* ... */ } 263 264 }; 265 266 // Deprecated, address_space_cast should be used instead. template <typename ElementType, access::address_space Space, access::decorated DecorateAddress> 267 multi_ptr<ElementType, Space, DecorateAddress> make_ptr(ElementType *); 268 269 270 template <access::address_space Space, access::decorated DecorateAddress,</pre> 271 typename ElementType> 272 multi_ptr<ElementType, Space, DecorateAddress> address_space_cast(ElementType *); 273 274 // Deduction guides 275 template <int dimensions, access::mode Mode, access::placeholder isPlaceholder,

| 276 | class T> |
|-----|--|
| 277 | multi_ptr(|
| 278 | <pre>accessor<t, access::target::global_buffer,="" dimensions,="" isplaceholder="" mode,="">)</t,></pre> |
| 279 | <pre>-> multi_ptr<t, access::address_space::global_space="">;</t,></pre> |
| 280 | <pre>template <int access::mode="" access::placeholder="" dimensions,="" isplaceholder,<="" mode,="" pre=""></int></pre> |
| 281 | class T> |
| 282 | <pre>multi_ptr(accessor<t, access::target::constant_buffer,<="" dimensions,="" mode,="" pre=""></t,></pre> |
| 283 | isPlaceholder>) |
| 284 | <pre>-> multi_ptr<t, access::address_space::constant_space="">;</t,></pre> |
| 285 | <pre>template <int access::mode="" access::placeholder="" dimensions,="" isplaceholder,<="" mode,="" pre=""></int></pre> |
| 286 | class T> |
| 287 | <pre>multi_ptr(accessor<t, access::target::local,="" dimensions,="" isplaceholder="" mode,="">)</t,></pre> |
| 288 | <pre>-> multi_ptr<t, access::address_space::local_space="">;</t,></pre> |
| 289 | |
| 290 | } // namespace sycl |

| Constructor | Description |
|--|--|
| <pre>template <typename access::<="" elementtype,="" pre=""></typename></pre> | Default constructor. |
| address_space Space, access::decorated | |
| DecorateAddress> | |
| <pre>multi_ptr()</pre> | |
| <pre>template <typename access::<="" elementtype,="" pre=""></typename></pre> | Copy constructor. |
| address_space Space, access::decorated | |
| DecorateAddress> | |
| <pre>multi_ptr(const multi_ptr &)</pre> | |
| <pre>template <typename access::<="" elementtype,="" pre=""></typename></pre> | Move constructor. |
| address_space Space, access::decorated | |
| DecorateAddress> | |
| <pre>multi_ptr(multi_ptr&&)</pre> | |
| <pre>template <typename access::<="" elementtype,="" pre=""></typename></pre> | Constructor that takes as an argument a dec- |
| address_space Space, access::decorated | orated pointer. |
| DecorateAddress> | |
| <pre>multi_ptr(multi_ptr<elementtype, space,="" yes="">::</elementtype,></pre> | |
| pointer) | |
| <pre>template <typename access::<="" elementtype,="" pre=""></typename></pre> | Constructor from a nullptr. |
| address_space Space, access::decorated | |
| DecorateAddress> | |
| <pre>multi_ptr(std::nullptr_t)</pre> | |
| <pre>template <typename access::<="" elementtype,="" pre=""></typename></pre> | Available only when: Space == access:: |
| <pre>address_space Space = access::address_space::</pre> | <pre>address_space::global_space.</pre> |
| global_space> | Constructs a multi_ptr <elementtype,< td=""></elementtype,<> |
| <pre>template <int access::mode="" dimensions,="" mode=""></int></pre> | <pre>access::address_space::global_space></pre> |
| multi_ptr(| from an accessor of access::target:: |
| <pre>accessor<elementtype, access::<="" dimensions,="" mode,="" pre=""></elementtype,></pre> | global_buffer. |
| <pre>target::global_buffer>)</pre> | |
| | Continued on next page |

Table 4.62: Constructors of the SYCL multi_ptr class template.

| Constructor | Description |
|--|---|
| <pre>template <typename access::<="" elementtype,="" pre=""></typename></pre> | Available only when: Space == access:: |
| <pre>address_space Space = access::address_space::</pre> | <pre>address_space::local_space.</pre> |
| local_space> | Constructs a <pre>multi_ptr<elementtype,< pre=""></elementtype,<></pre> |
| <pre>template <int access::mode="" dimensions,="" mode=""></int></pre> | <pre>access::address_space::local_space></pre> |
| multi_ptr(| from an accessor of <pre>access::target::</pre> |
| <pre>accessor<elementtype, access::<="" dimensions,="" mode,="" pre=""></elementtype,></pre> | local. |
| <pre>target::local>)</pre> | |
| <pre>template <typename access::<="" elementtype,="" pre=""></typename></pre> | Available only when: Space == access:: |
| <pre>address_space Space = access::address_space::</pre> | <pre>address_space::constant_space.</pre> |
| constant_space> | Constructs a <pre>multi_ptr<elementtype,< pre=""></elementtype,<></pre> |
| <pre>template <int access::mode="" dimensions,="" mode=""></int></pre> | <pre>access::address_space::constant_space</pre> |
| multi_ptr(| <pre>> from an accessor of access::target::</pre> |
| <pre>accessor<elementtype, access::<="" dimensions,="" mode,="" pre=""></elementtype,></pre> | constant_buffer. |
| <pre>target::constant_buffer>)</pre> | |
| <pre>template <typename access::<="" elementtype,="" pre=""></typename></pre> | Global function to create a multi_ptr in- |
| address_space Space, <pre>access::decorated</pre> | stance depending on the address space of |
| DecorateAddress> | the pointer type. An implementation must |
| <pre>multi_ptr<elementtype, decorateaddress="" space,=""></elementtype,></pre> | reject an argument if the deduced address |
| <pre>make_ptr(ElementType*)</pre> | space is not compatible with Space. |
| | End of table |

Table 4.62: Constructors of the SYCL multi_ptr class template.

| Operators | Description |
|--|--|
| <pre>template <typename access::address_space<="" pre="" value_type,=""></typename></pre> | Copy assignment operator. |
| <pre>Space, access::decorated DecorateAddress></pre> | |
| <pre>multi_ptr &operator=(const multi_ptr&)</pre> | |
| <pre>template <typename access::address_space<="" pre="" value_type,=""></typename></pre> | Move assignment operator. |
| <pre>Space, access::decorated DecorateAddress></pre> | |
| <pre>multi_ptr &operator=(multi_ptr&&)</pre> | |
| <pre>template <typename access::address_space<="" pre="" value_type,=""></typename></pre> | Assigns nullptr to the multi_ptr. |
| <pre>Space, access::decorated DecorateAddress></pre> | |
| <pre>multi_ptr &operator=(std::nullptr_t)</pre> | |
| <pre>template<access::address_space access::<="" asp,="" pre=""></access::address_space></pre> | Available only when: Space == access |
| decorated IsDecorated> | <pre>::address_space::generic_space &&</pre> |
| <pre>multi_ptr &operator=(const multi_ptr<value_type,< pre=""></value_type,<></pre> | ASP != access::address_space:: |
| ASP, IsDecorated>&) | constant_space. |
| | Assigns the value of the left hand side |
| | <pre>multi_ptr into the generic_ptr.</pre> |
| <pre>template<access::address_space access::<="" asp,="" pre=""></access::address_space></pre> | Available only when: Space == access |
| decorated IsDecorated> | <pre>::address_space::generic_space &&</pre> |
| <pre>multi_ptr &operator=(multi_ptr<value_type, asp,<="" pre=""></value_type,></pre> | ASP != access::address_space:: |
| IsDecorated>&& | constant_space. |
| | Move the value of the left hand side |
| | <pre>multi_ptr into the generic_ptr.</pre> |
| | Continued on next page |

| Table 4.63: | Operators | of multi_ | _ptr class. |
|-------------|-----------|-----------|-------------|
|-------------|-----------|-----------|-------------|

| Operators | Description |
|--|---|
| <pre>template <typename access::address_space<="" pre="" value_type,=""></typename></pre> | Available only when: !std::is_void< |
| <pre>Space, access::decorated DecorateAddress></pre> | value_type>::value. |
| <pre>pointer operator->()const</pre> | Returns the underlying pointer. |
| <pre>template <typename access::address_space<="" pre="" value_type,=""></typename></pre> | Available only when: !std::is_void< |
| <pre>Space, access::decorated DecorateAddress></pre> | value_type>::value. |
| reference operator*()const | Returns a reference to the pointed value. |
| <pre>template <typename access::address_space<="" pre="" value_type,=""></typename></pre> | Implicit conversion to the underlying pointer |
| <pre>Space, access::decorated DecorateAddress></pre> | type. Deprecated: The member function |
| operator pointer()const | get should be used instead |
| <pre>template <access::decorated isdecorated=""></access::decorated></pre> | Available only when: Space == access:: |
| <pre>operator multi_ptr<value_type, access::<="" pre=""></value_type,></pre> | <pre>address_space::generic_space.</pre> |
| <pre>address_space::private_space, IsDecorated>()const</pre> | Conversion from generic_ptr to |
| | private_ptr. The result is undefined |
| | if the pointer does not address the private |
| | address space. |
| <pre>template <access::decorated isdecorated=""></access::decorated></pre> | Available only when: Space == access:: |
| <pre>operator multi_ptr<const access::<="" pre="" value_type,=""></const></pre> | address_space::generic_space. |
| address_space::private_space, IsDecorated>()const | Conversion from generic_ptr to |
| | private_ptr. The result is undefined |
| | if the pointer does not address the private |
| | address space. |
| template <access::decorated isdecorated=""></access::decorated> | Available only when: Space == access:: |
| operator multi_ptr <value_type, access::<="" td=""><td>address_space::generic_space.</td></value_type,> | address_space::generic_space. |
| address_space::global_space, lsDecorated>()const | Conversion from generic_ptr to |
| | global_ptr. The result is undefined if |
| | address areas |
| tomplate cases udeconsted IsDeconsted | Available only when: Space = accoss t |
| operator multi ntr <const access:<="" td="" type="" value=""><td>Address space: generic space</td></const> | Address space: generic space |
| address space: global space IsDecorated () const | Conversion from generic ntr to |
| audress_spacegrobar_space, isbeebratear ()const | alobal ntr The result is undefined if |
| | the pointer does not address the global |
| | address space. |
| template <access::decorated isdecorated=""></access::decorated> | Available only when: Space == access:: |
| operator multi ptr <value access::<="" td="" type.=""><td>address space::generic space.</td></value> | address space::generic space. |
| address_space::local_space, IsDecorated>()const | Conversion from generic_ptr to |
| | local_ptr. The result is undefined if |
| | the pointer does not address the local |
| | address space. |
| <pre>template <access::decorated isdecorated=""></access::decorated></pre> | Available only when: Space == access:: |
| <pre>operator multi_ptr<const access::<="" pre="" value_type,=""></const></pre> | address_space::generic_space. |
| <pre>address_space::local_space, IsDecorated>()const</pre> | Conversion from generic_ptr to |
| | local_ptr. The result is undefined if |
| | the pointer does not address the local |
| | address space. |
| | Continued on next page |

Table 4.63: Operators of multi_ptr class.

| Operators | Description |
|---|--|
| <pre>template <access::decorated isdecorated=""></access::decorated></pre> | Available only when: !std::is_void< |
| <pre>operator multi_ptr<void, isdecorated="" space,="">()</void,></pre> | <pre>value_type>::value && !std::is_const<</pre> |
| const | value_type>::value. |
| | Implicit conversion to a multi_ptr of type |
| | void. |
| <pre>template <access::decorated isdecorated=""></access::decorated></pre> | Available only when: !std::is_void< |
| <pre>operator multi_ptr<const pre="" space,<="" void,=""></const></pre> | <pre>value_type>::value && std::is_const<</pre> |
| IsDecorated>()const | value_type>::value. |
| | Implicit conversion to a multi_ptr of type |
| | const void. |
| <pre>template <access::decorated isdecorated=""></access::decorated></pre> | Implicit conversion to a multi_ptr of type |
| <pre>operator multi_ptr<const pre="" space,<="" value_type,=""></const></pre> | <pre>const value_type.</pre> |
| IsDecorated>()const | |
| | Available only when: is_decorated == |
| <pre>operator multi_ptr<const pre="" space,<="" value_type,=""></const></pre> | true. |
| <pre>access::decorated::no>()const</pre> | Implicit conversion to the equivalent |
| | <pre>multi_ptr object that does not expose</pre> |
| | decorated pointers or references. |
| | Available only when: is_decorated == |
| <pre>operator multi_ptr<const pre="" space,<="" value_type,=""></const></pre> | false. |
| <pre>access::decorated::yes>()const</pre> | Implicit conversion to the equivalent |
| | <pre>multi_ptr object that exposes decorated</pre> |
| | pointers and references. |
| | End of table |

Table 4.63: Operators of multi_ptr class.

| Member function | Description |
|---|--|
| <pre>pointer get()const</pre> | Returns the underlying pointer. Whether the |
| | pointer is decorated depends on the value of |
| | DecorateAddress. |
| <pre>unspecified * get_decorated()const</pre> | Returns the underlying pointer decorated by |
| | the address space that it addressed. Note that |
| | the support involves implementation defined |
| | device compiler extensions. |
| <pre>std::add_pointer_t<value_type> get_raw()const</value_type></pre> | Returns the underlying pointer, always un- |
| | decorated. |
| <pre>void prefetch(size_t numElements)const</pre> | Available only when: Space == access:: |
| | address_space::global_space. |
| | Prefetches a number of elements specified |
| | by numElements into the global memory |
| | cache. This operation is an implementa- |
| | tion defined optimization and does not effect |
| | the functional behavior of the SYCL kernel |
| | function. |
| | End of table |

| Table 4.64: | Member | functions | of multi_ | _ptr class. |
|-------------|--------|-----------|-----------|-------------|
| | | | | |

| reference operator*(const multi_ptr& mp) Available only when: !std::is_void< ElementType>::value. Operator that returns a reference to the value_type of mp. multi_ptr& operator++(multi_ptr& mp) Available_only_when: _lstd::is_void |
|---|
| ElementType>::value. Operator that returns a reference to the value_type of mp. multi_ntr& operator++(multi_ntr& mp) Available_only_when:std::is_void |
| Operator that returns a reference to the value_type of mp. multi_ntr& operator++(multi_ntr& mp) |
| value_type of mp. multi ntr& operator++(multi ntr& mp) Available only when: |
| Available only when: Istd::is void |
| indici_perd operator (indici_perd mp) |
| <pre>ElementType>::value.</pre> |
| Increments mp by 1 and returns mp. |
| <pre>multi_ptr operator++(multi_ptr& mp, int) Available only when: !std::is_void<</pre> |
| <pre>ElementType>::value.</pre> |
| Increments mp by 1 and returns a new |
| multi_ptr with the value of the original mp. |
| <pre>multi_ptr& operator(multi_ptr& mp) Available only when: !std::is_void<</pre> |
| <pre>ElementType>::value.</pre> |
| Decrements mp by 1 and returns mp. |
| <pre>multi_ptr operator(multi_ptr& mp, int) Available only when: !std::is_void<</pre> |
| <pre>ElementType>::value.</pre> |
| Decrements mp by 1 and returns a new |
| multi_ptr with the value of the original mp. |
| <pre>multi_ptr& operator+=(multi_ptr& lhs,</pre> |
| difference_type r) ElementType>::value. |
| Moves mp forward by r and returns 1hs. |
| <pre>multi_ptr& operator-=(multi_ptr& lhs,</pre> |
| difference_type r) ElementType>::value. |
| Moves mp backward by r and returns 1hs. |
| <pre>multi_ptr operator+(const multi_ptr& lhs, Available only when: !std::is_void</pre> |
| difference_type r) ElementType>::value. |
| Creates a new multi_ptr that points r for- |
| ward compared to 1hs. |
| multi_ptr operator-(const multi_ptr& lhs, Available only when: !std::is_void Nicc Nicc |
| difference_type r) ElementType>::value. |
| Creates a new multi_ptr that points r back- |
| ward compared to Ins. |
| bool operator == (const multi_ptr& ins, const Comparison operator == for multi_ptr |
| Multi_ptr& rils) class. |
| wilti ntre rho) |
| heal energy (const multi ntre lbs, const multi ntre Comparison operator of for multi ntre class. |
| & rhc) |
| bool operators (const multi ntre lbc, const multi ntre Comparison operators for multi ntre class |
| & rbs) |
| bool operator <= (const multi ntr& lbs_const Comparison_operator <= for multi ntr |
| multi ntr& rhs) |
| bool operator>=(const multi ntr& lbs_const Comparison_operator >= for multi ntr |
| multi ntr& rhc) |
| bool operator (const multi ntr& lbs_std: mullintr t Comparison operator for multi ntr class. |
|) with a std::mullintr + |
| Continued on next page |

Table 4.65: Hidden friend functions of the multi_ptr class.

| Hidden friend function | Description |
|--|--|
| <pre>bool operator!=(const multi_ptr& lhs, std::nullptr_t</pre> | Comparison operator != for multi_ptr class |
|) | with a std::nullptr_t. |
| <pre>bool operator<(const multi_ptr& lhs, std::nullptr_t)</pre> | Comparison operator < for multi_ptr class |
| | with a std::nullptr_t. |
| <pre>bool operator>(const multi_ptr& lhs, std::nullptr_t)</pre> | Comparison operator > for multi_ptr class |
| | with a std::nullptr_t. |
| <pre>bool operator<=(const multi_ptr& lhs, std::nullptr_t</pre> | Comparison operator <= for multi_ptr class |
|) | with a std::nullptr_t. |
| <pre>bool operator>=(const multi_ptr& lhs, std::nullptr_t</pre> | Comparison operator >= for multi_ptr class |
|) | with a std::nullptr_t. |
| <pre>bool operator==(std::nullptr_t, const multi_ptr& rhs</pre> | Comparison operator == for multi_ptr class |
|) | with a std::nullptr_t. |
| <pre>bool operator!=(std::nullptr_t, const multi_ptr& rhs</pre> | Comparison operator != for multi_ptr class |
|) | with a std::nullptr_t. |
| <pre>bool operator<(std::nullptr_t, const multi_ptr& rhs)</pre> | Comparison operator < for multi_ptr class |
| | with a std::nullptr_t. |
| <pre>bool operator>(std::nullptr_t, const multi_ptr& rhs)</pre> | Comparison operator > for multi_ptr class |
| | with a std::nullptr_t. |
| <pre>bool operator<=(std::nullptr_t, const multi_ptr& rhs</pre> | Comparison operator <= for multi_ptr class |
|) | with a std::nullptr_t. |
| <pre>bool operator>=(std::nullptr_t, const multi_ptr& rhs</pre> | Comparison operator >= for multi_ptr class |
|) | with a std::nullptr_t. |
| | End of table |

Table 4.65: Hidden friend functions of the multi_ptr class.

The following is the overview of the legacy interface from 1.2.1 provided for the multi_ptr class.

```
1 namespace sycl {
2
3 // Legacy interface, inherited from 1.2.1.
   // Deprecated.
4
   template <typename ElementType, access::address_space Space>
5
   class [[deprecated]] multi_ptr<ElementType, Space, access::decorated::legacy> {
6
7
    public:
8
     using element_type = ElementType;
9
     using difference_type = std::ptrdiff_t;
10
     // Implementation defined pointer and reference types that correspond to
11
12
     // SYCL/OpenCL interoperability types for OpenCL C functions.
13
     using pointer_t = multi_ptr<ElementType, Space, access::decorated::yes>::pointer;
14
     using const_pointer_t = multi_ptr<const ElementType, Space, access::decorated::yes>::pointer;
15
     using reference_t = multi_ptr<ElementType, Space, access::decorated::yes>::reference;
16
     using const_reference_t = multi_ptr<const ElementType, Space, access::decorated::yes>::reference
          ;
17
18
     static constexpr access::address_space address_space = Space;
19
20
      // Constructors
21
     multi_ptr();
```

```
22
      multi_ptr(const multi_ptr&);
23
      multi_ptr(multi_ptr&&);
24
      multi_ptr(pointer_t);
25
      multi_ptr(ElementType*);
26
      multi_ptr(std::nullptr_t);
27
      ~multi_ptr();
28
29
      // Assignment and access operators
30
      multi_ptr &operator=(const multi_ptr&);
31
      multi_ptr &operator=(multi_ptr&&);
32
      multi_ptr &operator=(pointer_t);
33
      multi_ptr &operator=(ElementType*);
34
      multi_ptr &operator=(std::nullptr_t);
35
      friend ElementType& operator*(const multi_ptr& mp) { /* ... */ }
36
      ElementType* operator->() const;
37
38
      // Only if Space == global_space
39
      template <int dimensions, access::mode Mode, access::placeholder isPlaceholder>
40
      multi_ptr(accessor<ElementType, dimensions, Mode, access::target::global_buffer, isPlaceholder>)
          ;
41
42
      // Only if Space == local_space
43
      template <int dimensions, access::mode Mode, access::placeholder isPlaceholder>
44
      multi_ptr(accessor<ElementType, dimensions, Mode, access::target::local, isPlaceholder>);
45
46
      // Only if Space == constant_space
47
      template <int dimensions, access::mode Mode, access::placeholder isPlaceholder>
48
      multi_ptr(accessor<ElementType, dimensions, Mode, access::target::constant_buffer, isPlaceholder</pre>
          >);
49
50
      // Returns the underlying OpenCL C pointer
51
      pointer_t get() const;
52
53
      // Implicit conversion to the underlying pointer type
54
      operator ElementType*() const;
55
56
      // Implicit conversion to a multi_ptr<void>
57
      // Only available when ElementType is not const-qualified
58
      operator multi_ptr<void, Space>() const;
59
60
      // Implicit conversion to a multi_ptr<const void>
61
      // Only available when ElementType is const-qualified
62
      operator multi_ptr<const void, Space>() const;
63
64
      // Implicit conversion to multi_ptr<const ElementType, Space>
65
      operator multi_ptr<const ElementType, Space>() const;
66
67
      // Arithmetic operators
68
      friend multi_ptr& operator++(multi_ptr& mp) { /* ... */ }
69
      friend multi_ptr operator++(multi_ptr& mp, int) { /* ... */ }
70
      friend multi_ptr& operator--(multi_ptr& mp) { /* ... */ }
71
      friend multi_ptr operator--(multi_ptr& mp, int) { /* ... */ }
72
      friend multi_ptr& operator+=(multi_ptr& lhs, difference_type r) { /* ... */ }
73
      friend multi_ptr& operator-=(multi_ptr& lhs, difference_type r) { /* ... */ }
74
      friend multi_ptr operator+(const multi_ptr& lhs, difference_type r) { /* ... */ }
```

```
friend multi_ptr operator-(const multi_ptr& lhs, difference_type r) { /* ... */ }
 75
 76
 77
       void prefetch(size_t numElements) const;
 78
 79
       friend bool operator==(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }
 80
       friend bool operator!=(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }
 81
       friend bool operator<(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }</pre>
 82
       friend bool operator>(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }
 83
       friend bool operator<=(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }
 84
       friend bool operator>=(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }
 85
 86
       friend bool operator==(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }
       friend bool operator!=(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }
 87
 88
       friend bool operator<(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }</pre>
       friend bool operator>(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }
 89
 90
       friend bool operator<=(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }</pre>
 91
       friend bool operator>=(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }
 92
 93
       friend bool operator==(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }
 94
       friend bool operator!=(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }
 95
       friend bool operator<(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }</pre>
 96
       friend bool operator>(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }
 97
       friend bool operator<=(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }</pre>
 98
       friend bool operator>=(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }
99
100 };
101
102 // Legacy interface, inherited from 1.2.1.
103 // Deprecated.
104 // Specialization of multi_ptr for void and const void
105 // VoidType can be either void or const void
106
    template <access::address_space Space>
107
     class [[deprecated]] multi_ptr<VoidType, Space, access::decorated::legacy> {
108
     public:
109
      using element_type = VoidType;
110
       using difference_type = std::ptrdiff_t;
111
112
       // Implementation defined pointer types that correspond to
113
       // SYCL/OpenCL interoperability types for OpenCL C functions
114
       using pointer_t = multi_ptr<VoidType, Space, access::decorated::yes>::pointer;
115
       using const_pointer_t = multi_ptr<const VoidType, Space, access::decorated::yes>::pointer;
116
117
       static constexpr access::address_space address_space = Space;
118
119
       // Constructors
120
      multi_ptr();
121
      multi_ptr(const multi_ptr&);
122
      multi_ptr(multi_ptr&&);
123
      multi_ptr(pointer_t);
124
      multi_ptr(VoidType*);
125
      multi_ptr(std::nullptr_t);
126
       ~multi_ptr();
127
128
       // Assignment operators
129
       multi_ptr &operator=(const multi_ptr&);
```

CHAPTER 4. SYCL PROGRAMMING INTERFACE

191

```
130
       multi_ptr &operator=(multi_ptr&&);
131
       multi_ptr &operator=(pointer_t);
132
       multi_ptr &operator=(VoidType*);
133
       multi_ptr &operator=(std::nullptr_t);
134
135
       // Only if Space == global_space
136
       template <typename ElementType, int dimensions, access::mode Mode>
137
       multi_ptr(accessor<ElementType, dimensions, Mode, access::target::global_buffer>);
138
139
       // Only if Space == local_space
140
       template <typename ElementType, int dimensions, access::mode Mode>
141
       multi_ptr(accessor<ElementType, dimensions, Mode, access::target::local>);
142
143
       // Only if Space == constant_space
144
       template <typename ElementType, int dimensions, access::mode Mode>
145
       multi_ptr(accessor<ElementType, dimensions, Mode, access::target::constant_buffer>);
146
147
       // Returns the underlying OpenCL C pointer
148
       pointer_t get() const;
149
150
       // Implicit conversion to the underlying pointer type
151
       operator VoidType*() const;
152
153
       // Explicit conversion to a multi_ptr<ElementType>
154
       // If VoidType is const, ElementType must be as well
155
       template <typename ElementType>
156
       explicit operator multi_ptr<ElementType, Space>() const;
157
158
       // Implicit conversion to multi_ptr<const void, Space>
159
       operator multi_ptr<const void, Space>() const;
160
161
       friend bool operator==(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }
162
       friend bool operator!=(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }
163
       friend bool operator<(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }</pre>
       friend bool operator>(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }
164
165
       friend bool operator<=(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }</pre>
166
       friend bool operator>=(const multi_ptr& lhs, const multi_ptr& rhs) { /* ... */ }
167
       friend bool operator==(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }
168
169
       friend bool operator!=(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }
170
       friend bool operator<(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }</pre>
171
       friend bool operator>(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }
172
       friend bool operator<=(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }</pre>
173
       friend bool operator>=(const multi_ptr& lhs, std::nullptr_t) { /* ... */ }
174
175
       friend bool operator==(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }
       friend bool operator!=(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }
176
       friend bool operator<(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }</pre>
177
178
       friend bool operator>(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }
179
       friend bool operator<=(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }</pre>
       friend bool operator>=(std::nullptr_t, const multi_ptr& rhs) { /* ... */ }
180
181
182 };
183
184 } // namespace sycl
```

4.7.7.2 Explicit pointer aliases

SYCL provides aliases to the multi_ptr class template (see Section 4.7.7.1) for each specialization of access:: address_space.

A synopsis of the SYCL multi_ptr class template aliases is provided below.

```
1 namespace sycl {
2
3
   template <typename ElementType, access::address_space Space, access::decorated IsDecorated>
4
   class multi_ptr;
5
6
    // Template specialization aliases for different pointer address spaces
7
8
   template <typename ElementType, access::decorated IsDecorated = access::decorated::legacy>
9
    using global_ptr = multi_ptr<ElementType, access::address_space::global_space,</pre>
10
                                 IsDecorated>;
11
12 template <typename ElementType, access::decorated IsDecorated = access::decorated::legacy>
13
   using local_ptr = multi_ptr<ElementType, access::address_space::local_space,</pre>
14
                                 IsDecorated>;
15
16
   template <typename ElementType, access::decorated IsDecorated = access::decorated::legacy>
17
    using constant_ptr = multi_ptr<ElementType, access::address_space::constant_space,</pre>
18
                                   IsDecorated>;
19
20
   template <typename ElementType, access::decorated IsDecorated = access::decorated::legacy>
21
    using private_ptr = multi_ptr<ElementType, access::address_space::private_space,</pre>
22
                                   IsDecorated>;
23
24
   // Template specialization aliases for different pointer address spaces.
25
   // The interface exposes non-decorated pointer while keeping the
26
   // address space information internally.
27
28
   template <typename ElementType>
29
   using raw_global_ptr = multi_ptr<ElementType, access::address_space::global_space,</pre>
30
                                      access::decorated::no>;
31
32
   template <typename ElementType>
33
    using raw_local_ptr = multi_ptr<ElementType, access::address_space::local_space,</pre>
34
                                     access::decorated::no>;
35
36
   template <typename ElementType>
37
    using raw_constant_ptr = multi_ptr<ElementType, access::address_space::constant_space,</pre>
38
                                        access::decorated::no>;
39
40 template <typename ElementType>
41
    using raw_private_ptr = multi_ptr<ElementType, access::address_space::private_space,</pre>
42
                                       access::decorated::no>;
43
44
   // Template specialization aliases for different pointer address spaces.
45
   // The interface exposes decorated pointer.
46
47 template <typename ElementType>
   using decorated_global_ptr = multi_ptr<ElementType, access::address_space::global_space,</pre>
48
```

```
49
                                           access::decorated::yes>;
50
51
    template <typename ElementType>
52
    using decorated_local_ptr = multi_ptr<ElementType, access::address_space::local_space,</pre>
53
                                          access::decorated::yes>;
54
55 template <typename ElementType>
56 using decorated_constant_ptr = multi_ptr<ElementType, access::address_space::constant_space,
57
                                             access::decorated::yes>;
58
59 template <typename ElementType>
60 using decorated_private_ptr = multi_ptr<ElementType, access::address_space::private_space,
61
                                            access::decorated::yes>;
62
63 } // namespace sycl
```

Note that using global_ptr, local_ptr, constant_ptr or private_ptr without specifying the decoration is deprecated. The default argument is provided for compatibility with 1.2.1.

4.7.8 Samplers

The SYCL sampler struct encapsulates a configuration for sampling a sampled_image. A SYCL sampler can map to one native backend object.

```
1 namespace sycl {
2 enum class addressing_mode: unsigned int {
3
     mirrored_repeat,
4
     repeat,
 5
     clamp_to_edge,
 6
     clamp,
7
     none
8 };
9
10 enum class filtering_mode: unsigned int {
11
     nearest.
12
     linear
13 };
14
15 enum class coordinate_normalization_mode : unsigned int {
     normalized,
16
     unnormalized
17
18 };
19
20 struct image_sampler {
      addressing_mode addressing;
21
22
     coordinate_mode coordinate;
23
     filtering_mode filtering;
24 };
25 } // namespace sycl
```

| addressing_mode | Description |
|-----------------|--|
| mirrored_repeat | Out of range coordinates will be flipped |
| | at every integer junction. This addressing |
| | mode can only be used with normalized co- |
| | ordinates. If normalized coordinates are not |
| | used, this addressing mode may generate |
| | image coordinates that are undefined. |
| repeat | Out of range image coordinates are wrapped |
| | to the valid range. This addressing mode can |
| | only be used with normalized coordinates. |
| | If normalized coordinates are not used, this |
| | addressing mode may generate image coor- |
| | dinates that are undefined. |
| clamp_to_edge | Out of range image coordinates are clamped |
| | to the extent. |
| clamp | Out of range image coordinates will return |
| | a border color. |
| none | For this addressing mode the programmer |
| | guarantees that the image coordinates used |
| | to sample elements of the image refer to a |
| | location inside the image; otherwise the re- |
| | sults are undefined. |
| | End of table |

Table 4.66: Addressing modes description.

| filtering_mode | Description |
|----------------|---|
| nearest | Chooses a color of nearest pixel. |
| linear | Performs a linear sampling of adjacent pix- |
| | els. |
| | End of table |

Table 4.67: Filtering modes description.

| coordinate_normalization_mode | Description |
|-------------------------------|---------------------------------------|
| normalized | Normalizes image coordinates. |
| unnormalized | Does not normalize image coordinates. |
| | End of table |

Table 4.68: Coordinate normalization modes description.

| | T |
|---|---|
| Constructor | Description |
| sampler(| Constructs a SYCL sampler instance with |
| <pre>coordinate_normalization_mode normalizationMode,</pre> | address mode, filtering mode and coordinate |
| addressing_mode addressingMode, | normalization mode specified by the respec- |
| <pre>filtering_mode filteringMode,</pre> | tive parameters. It is not valid to construct |
| <pre>const property_list &propList = {})</pre> | a SYCL sampler within a SYCL kernel |
| | function. The optional property_list pro- |
| | vides properties for the constructed SYCL |
| | sampler object. |
| | End of table |

| Table 4.69: | Constructors | the | sampler | class. |
|-------------|--------------|-----|---------|--------|
|-------------|--------------|-----|---------|--------|

| Member function | Description |
|---|---|
| <pre>addressing_mode get_addressing_mode()const</pre> | Return the addressing mode used to con- |
| | struct this SYCL sampler. |
| <pre>filtering_mode get_filtering_mode()const</pre> | Return the filtering mode used to construct |
| | this SYCL sampler. |
| coordinate_normalization_mode | Return the coordinate normalization mode |
| <pre>get_coordinate_normalization_mode()const</pre> | used to construct this SYCL sampler. |
| | End of table |

Table 4.70: Member functions for the sampler class.

4.8 Unified shared memory (USM)

This section describes new properties and routines for pointer-based memory management interfaces in SYCL. These routines augment, rather than replace, the existing buffer-based interfaces in SYCL.

Unified Shared Memory (USM) provides a pointer-based alternative to the buffer programming model. USM enables:

- Easier integration into existing code bases by representing allocations as pointers rather than buffers, with full support for pointer arithmetic into allocations.
- Fine-grain control over ownership and accessibility of allocations, to optimally choose between performance and programmer convenience.
- A simpler programming model, by automatically migrating some allocations between SYCL devices and the host.

4.8.1 Unified addressing

Unified Addressing guarantees that all devices will use a unified address space. Pointer values in the unified address space will always refer to the same location in memory. The unified address space encompasses the host and one or more devices. Note that this does not require addresses in the unified address space to be accessible on all devices, just that pointer values will be consistent.

4.8.2 Kinds of unified shared memory

USM is a capability that, when available, provides the ability to create allocations that are visible to both host and device(s). USM builds upon Unified Addressing to define a shared address space where pointer values in this space always refer to the same location in memory. USM defines multiple tiers of increasing capability described in the following sections:

- Explicit USM
- Restricted USM
- Concurrent USM
- System USM

USM is an optional feature which may not be supported by all devices, and devices that support USM may only support some of these tiers. A SYCL application can use the device::has() function to determine the level of USM support for a device. See Table 4.20 in Section 4.6.4.3 for more details.

4.8.2.1 Explicit USM

Explicit USM defines capabilities for explicitly managing device memory. Programmers directly allocate device memory, and data must be explicitly copied between the host and a device. Device allocations are obtained through SYCL USM device allocation routines instead of system allocation routines like std::malloc or C++ new. Device allocations are not accessible on the host, but the pointer values remain consistent on account of Unified Addressing. Greater detail about how allocations are used is described by the following tables.

4.8.2.2 Restricted USM

Restricted USM defines capabilities for implicitly sharing data between host and devices. However, Restricted USM, as the name implies, is limited in that host and device may not concurrently compute on memory in the shared address space. Restricted USM builds upon Explicit USM by adding two new types of allocations, host and shared. Allocations are obtained through SYCL allocator instead of the system allocator. shared allocations may be limited by device memory. Greater detail about the allocation types defined in Restricted USM and their usage is described by the following tables.

4.8.2.3 Concurrent USM

Concurrent USM builds upon Restricted USM by enabling concurrent access to shared allocations between host and devices. Additionally, some implementations may support a working set of shared allocations larger than device memory.

4.8.2.4 System USM

System USM extends upon the previous tiers by performing all shared allocations with the normal system memory allocation routines. In particular, programmers may now use std::malloc or C++ new instead of USM allocation routines to create shared allocations. Likewise, std::free and delete are used instead of sycl::free. Note that host and device allocations are unaffected by this change and must still be allocated using their respective USM functions in order to guarantee their behavior.

| USM allocation type | Description |
|---------------------|--|
| host | Allocations in host memory that are accessible by a device |
| device | Allocations in device memory that are not accessible by the host |
| shared | Allocations in shared memory that are accessible by both host and device |

| Allocation Type | Initial Location | Accessible By | | Migratable To | |
|-----------------|------------------|----------------|----------------|----------------|----------|
| device | device | host | No | host | No |
| | | device | Yes | device | N/A |
| | | Another device | Optional (P2P) | Another device | No |
| host | host | host | Yes | host | N/A |
| | | Any device | Yes | device | No |
| | | host | Yes | host | Yes |
| shared | Unspecified | device | Yes | device | Yes |
| | | Another device | Optional (P2P) | Another device | Optional |

Table 4.71: Type of USM allocations.

Table 4.72: Characteristics of the different kinds of USM allocation.

4.8.3 USM allocations

There are different types of allocation described on Table 4.71 with different characteristics exposed on Table 4.72 which can be referred using the following enum:

```
1 namespace sycl {
2
     namespace usm {
3
       enum class alloc {
4
         host,
5
          device,
6
          shared,
7
          unknown
8
       };
9
      }
10 }
```

4.8.4 C++ allocator interface

```
1 template <typename T, usm::alloc AllocKind, size_t Alignment = 0>
 2 class usm_allocator {
 3
    public:
 4
      using value_type = T;
 5
 6
    public:
 7
      template <typename U> struct rebind {
 8
        typedef usm_allocator<U, AllocKind, Alignment> other;
9
      };
10
11
      usm_allocator() noexcept = delete;
12
      usm_allocator(const context &ctxt, const device &dev) noexcept;
13
      usm_allocator(const queue &q) noexcept;
```

14 usm_allocator(const usm_allocator &other) noexcept; 15 16 template <class U> usm_allocator(usm_allocator<U, AllocKind, Alignment> const &) noexcept; 17 18 /// Allocate memory 19 T *allocate(size_t Size); 20 21 /// Deallocate memory void deallocate(T *Ptr, size_t size); 22 23 24 /// Constructs an object on memory pointed by Ptr. 25 111 26 /// Note: AllocKind == alloc::device is not allowed. 27 template <</pre> 28 usm::alloc AllocT = AllocKind, 29 typename std::enable_if<AllocT != usm::alloc::device, int>::type = 0, 30 class U, class... ArgTs> 31 void construct(U *Ptr, ArgTs &&... Args); 32 33 /// Throws an error when trying to construct a device allocation 34 /// on the host 35 template <</pre> usm::alloc AllocT = AllocKind, 36 37 typename std::enable_if<AllocT == usm::alloc::device, int>::type = 0, class U, class... ArgTs> 38 39 void construct(U *Ptr, ArgTs &&... Args); 40 41 /// Destroys an object. 42 111 /// Note:: AllocKind == alloc::device is not allowed 43 44 template <</pre> 45 usm::alloc AllocT = AllocKind, 46 typename std::enable_if<AllocT != usm::alloc::device, int>::type = 0> 47 void destroy(T *Ptr); 48 49 /// Throws an error when trying to destroy a device allocation 50 /// on the host 51 template <</pre> 52 usm::alloc AllocT = AllocKind, 53 typename std::enable_if<AllocT == usm::alloc::device, int>::type = 0> 54 void destroy(T *Ptr); 55 }; 56 57 /// Equality Comparison 58 /// 59 /// Allocators only compare equal if they are of the same USM kind, alignment, 60 /// context, and device (when kind is not host) 61 template <class T, usm::alloc AllocKindT, size_t AlignmentT, class U, 62 usm::alloc AllocKindU, size_t AlignmentU> 63 bool operator==(const usm_allocator<T, AllocKindT, AlignmentT> &, const usm_allocator<U, AllocKindU, AlignmentU> &) noexcept; 64 65 66 /// Inequality Comparison 67 /// 68 /// Allocators only compare unequal if they are not of the same USM kind, alignment,

```
69 /// context, or device (when kind is not host)
70 template <class T, class U, usm::alloc AllocKind, size_t Alignment = 0>
71 bool operator!=(const usm_allocator<T, AllocKind, Alignment> &allocT,
72 const usm_allocator<U, AllocKind, Alignment> &allocU) noexcept;
```

4.8.5 Utility functions

While the modern C++ usm_allocator interface is sufficient for specifying USM allocations and deallocations, many programmers may prefer C-style malloc-influenced APIs. As a convenience to programmers, malloc-style APIs are also defined. Additionally, other utility functions are specified in the following sections to perform various operations such as memory copies and initializations as well as to provide performance hints.

4.8.5.1 Explicit USM

```
4.8.5.1.1 malloc_device
```

```
1
   (1)
2 void* sycl::malloc_device(size_t num_bytes,
3
                             const sycl::device& dev,
4
                              const sycl::context& ctxt);
5
6
   (2)
7
   template <typename T>
8
   T* sycl::malloc_device(size_t count,
0
                           const sycl::device& dev,
10
                           const sycl::context& ctxt);
```

Parameters

- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::device& dev the SYCL device to allocate on
- const sycl::context& ctxt the SYCL context to which device belongs

Return value Returns a pointer to the newly allocated memory on the specified device on success. This memory

is not accessible on the host. Memory allocated by sycl::malloc_device must be deallocated with sycl ::free to avoid memory leaks. If ctxt is a host context, it should behave as if calling malloc_host. On failure, returns nullptr.

Parameters

CHAPTER 4. SYCL PROGRAMMING INTERFACE

200

- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::queue& q the SYCL q that provides the device and context to allocate against
- **Return value** Returns a pointer to the newly allocated memory on the device associated with q on success. This memory is not accessible on the host. Memory allocated by sycl::malloc_device must be deallocated with sycl::free to avoid memory leaks. If q is a host queue, it should behave as if calling malloc_host. On failure, returns nullptr.

4.8.5.1.2 aligned_alloc_device

```
1 (1)
    void* sycl::aligned_alloc_device(size_t alignment,
2
3
                                     size_t num_bytes,
4
                                     const sycl::device& dev,
5
                                     const sycl::context& ctxt);
6
7
    (2)
8
   template <typename T>
9
   T* sycl::aligned_alloc_device(size_t alignment,
10
                                  size_t count,
11
                                  const sycl::device& dev,
12
                                  const sycl::context& ctxt);
```

Parameters

- size_t alignment specifies the byte alignment. Must be a valid alignment supported by the implementation.
- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::device& dev the device to allocate on
- const sycl::context& ctxt the SYCL context to which device belongs
- **Return value** Returns a pointer to the newly allocated memory on the specified device on success. This memory is not accessible on the host. Memory allocated by sycl::aligned_alloc_device must be deallocated with sycl::free to avoid memory leaks. If ctxt is a host context, it should behave as if calling aligned_alloc_host. On failure, returns nullptr.

```
1 (1)
2
   void* sycl::aligned_alloc_device(size_t alignment,
3
                                     size_t size,
4
                                     const sycl::queue& q);
5
6
   (2)
7
   template <typename T>
8
   T* sycl::aligned_alloc_device(size_t alignment,
9
                                  size_t count,
```

10

const sycl::queue& q);

Parameters

- size_t alignment specifies the byte alignment. Must be a valid alignment supported by the implementation.
- (1) **size_t size** number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::queue& q the SYCL q that provides the device and context to allocate against
- **Return value** Returns a pointer to the newly allocated memory on the device associated with q on success. This memory is not accessible on the host. Memory allocated by sycl::aligned_alloc_device must be deallocated with sycl::free to avoid memory leaks. If q is a host queue, it should behave as if calling aligned_alloc_host. On failure, returns nullptr.

4.8.5.1.3 memcpy

```
1 class handler {
 2
     . . .
 3
     public:
 4
      . . .
 5
      void memcpy(void* dest, const void* src, size_t num_bytes);
 6 };
 7
 8
    class queue {
 9
     . . .
10
    public:
11
      . . .
12
      event memcpy(void* dest, const void* src, size_t num_bytes);
13 };
```

Parameters

- void* dest pointer to the destination memory
- const void* src pointer to the source memory
- size_t num_bytes number of bytes to copy

Return value Returns an event representing the copy operation.

4.8.5.1.4 memset

202

```
1 class handler {
2 ...
3 public:
4 ...
5 void memset(void* ptr, int value, size_t num_bytes);
6 };
```

```
7
8 class queue {
9 ...
10 public:
11 ...
12 event memset(void* ptr, int value, size_t num_bytes);
13 };
```

Parameters

- void* ptr pointer to the memory to fill
- int value value to be set. Value is interpreted as an unsigned char
- size_t num_bytes number of bytes to fill

Return value Returns an event representing the fill operation.

4.8.5.1.5 fill

```
1 class handler {
2
    . . .
3
    public:
4
     . . .
     template <typename T>
5
6
     void fill(void* ptr, const T& pattern, size_t count)
7
   };
8
9 class queue {
10
    . . .
11
   public:
12
     . . .
13
    template <typename T>
     event fill(void* ptr, const T& pattern, size_t count);
14
15 };
```

Parameters

- void* ptr pointer to the memory to fill
- const T& pattern pattern to be filled. T should be trivially copyable.
- size_t count number of times to fill pattern into ptr

Return value Returns an event representing the fill operation or void if on the handler.

4.8.5.2 Restricted USM

Restricted USM includes all of the Utility Functions of Explicit USM. It additionally introduces new functions to support **host** and **shared** allocations.

4.8.5.2.1 malloc

1 (1)

```
2 void* sycl::malloc_host(size_t num_bytes, const sycl::context& ctxt);
```

```
3 (2)
```

```
4 template <typename T>
```

```
5 T* sycl::malloc_host(size_t count, const sycl::context& ctxt);
```

Parameters

- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::context& ctxt the SYCL context that contains the devices that will access the host allocation
- **Return value** Returns a pointer to the newly allocated host memory on success. Memory allocated by sycl:: malloc_host must be deallocated with sycl::free to avoid memory leaks. On failure, returns nullptr.

```
1 (1)
2 void* sycl::malloc_host(size_t num_bytes, const sycl::queue& q);
3 (2)
4 template <typename T>
5 T* sycl::malloc_host(size_t count, const sycl::queue& q);
```

Parameters

- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::queue& q the SYCL queue whose context contains the devices that will access the host allocation

Return value Returns a pointer to the newly allocated host memory on success. Memory allocated by sycl:: malloc_host must be deallocated with sycl::free to avoid memory leaks. On failure, returns nullptr.

```
1 (1)
2
   void* sycl::malloc_shared(size_t num_bytes,
3
                             const sycl::device& dev,
4
                             const sycl::context& ctxt);
5
  (2)
6 template <typename T>
7
   T* sycl::malloc_shared(size_t count,
8
                          const sycl::device& dev,
9
                          const sycl::context& ctxt);
```

Parameters

- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::device& dev the SYCL device which shares access to this memory with the host. The

requested memory is allocated on this device.

- const sycl::context& ctxt defines the set of devices that may also share access to this memory, if supported. The dev device must be contained within this context.
- **Return value** Returns a pointer to the newly allocated shared memory on the specified device on success. Memory allocated by sycl::malloc_shared must be deallocated with sycl::free to avoid memory leaks. If ctxt is a host context, should behave as if calling malloc_host. On failure, returns nullptr.

Parameters

- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::queue& q the SYCL q that provides the device and context to allocate against
- **Return value** Returns a pointer to the newly allocated shared memory on the device associated with q on success. Memory allocated by sycl::malloc_shared must be deallocated with sycl::free to avoid memory leaks. On failure, returns nullptr.

4.8.5.2.2 aligned_alloc_host

```
1 (1)
```

- 2 void* sycl::aligned_alloc_host(size_t alignment, size_t num_bytes, const sycl::context& ctxt);
- 3 (2)
- 4 template <typename T>
- 5 T* sycl::aligned_alloc_host(size_t alignment, size_t count, const sycl::context& ctxt);

Parameters

- size_t alignment specifies the byte alignment. Must be a valid alignment supported by the implementation.
- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::context& ctxt the SYCL context that contains the devices that will access the host allocation
- **Return value** Returns a pointer to the newly allocated host memory on success. Memory allocated by sycl ::aligned_alloc_host must be deallocated with sycl::free to avoid memory leaks. On failure, returns nullptr.

1 (1)

```
2 void* sycl::aligned_alloc_host(size_t alignment, size_t num_bytes, const sycl::queue& q);
```

```
3 (2)
```

```
4 template <typename T>
```

```
5 void* sycl::aligned_alloc_host(size_t alignment, size_t count, const sycl::queue& q);
```

Parameters

- size_t alignment specifies the byte alignment. Must be a valid alignment supported by the implementation.
- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::queue& q the SYCL q whose context contains the devices that will access the host allocation

```
1 (1)
2 void* sycl::aligned_alloc_shared(size_t alignment,
3
                                    size_t num_bytes,
4
                                     const sycl::device& dev,
5
                                     const sycl::context& ctxt);
6 (2)
7 template <typename T>
8 T* sycl::aligned_alloc_shared(size_t alignment,
9
                                 size_t count,
                                 const sycl::device& dev,
10
11
                                 const sycl::context& ctxt);
```

Parameters

- size_t alignment specifies the byte alignment. Must be a valid alignment supported by the implementation.
- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::device& dev the SYCL device which shares access to this memory with the host. The requested memory is allocated on this device.
- const sycl::context& ctxt defines the set of devices that may also share access to this memory, if supported. The dev device must be contained within this context.
- **Return value** Returns a pointer to the newly allocated shared memory on the specified device on success. Memory allocated by sycl::aligned_alloc_shared must be deallocated with sycl::free to avoid memory leaks. If ctxt is a host context, should behave as if calling aligned_alloc_host. On failure, returns nullptr.

Return value Returns a pointer to the newly allocated host memory on success. Memory allocated by sycl ::aligned_alloc_host must be deallocated with sycl::free to avoid memory leaks. On failure, returns nullptr.

```
(1)
1
   void* sycl::aligned_alloc_shared(size_t alignment,
2
3
                                     size_t num_bytes,
4
                                     const sycl::queue& q);
5
   (2)
6
   template <typename T>
7
  T* sycl::aligned_alloc_shared(size_t alignment,
8
                                  size_t count,
9
                                  const sycl::queue& q);
```

Parameters

- size_t alignment specifies the byte alignment. Must be a valid alignment supported by the implementation.
- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::queue& q the SYCL q that provides the device and context to allocate against
- **Return value** Returns a pointer to the newly allocated shared memory on the device associated with q on success. Memory allocated by sycl::aligned_alloc_shared must be deallocated with sycl::free to avoid memory leaks. If ctxt is a host context, should behave as if calling aligned_alloc_host. On failure, returns nullptr.

4.8.5.2.3 Performance hints

Programmers may provide hints to the runtime that data should be made available on a device earlier than Unified Shared Memory would normally require it to be available. This can be accomplished through enqueueing prefetch commands. Prefetch commands may not be overlapped with kernel execution in Restricted USM.

4.8.5.2.3.1 prefetch

```
class handler {
1
2
     . . .
3
     public:
4
      . . .
5
      void prefetch(const void* ptr, size_t num_bytes);
6
    };
7
8
   class queue {
9
    . . .
10
    public:
11
12
      event prefetch(const void* ptr, size_t num_bytes);
13 };
```

Parameters

- const void* ptr pointer to the memory to be prefetched to the device
- size_t num_bytes number of bytes requested to be prefetched

Return value Returns an event representing the prefetch operation.

4.8.5.3 Concurrent USM

Concurrent USM contains all the utility functions of Explicit USM and Restricted USM. It introduces a new function, sycl::queue::mem_advise, that allows programmers to provide additional information to the underlying runtime about how different allocations are used.

4.8.5.3.1 Performance hints

4.8.5.3.1.1 prefetch

In Concurrent USM, prefetch commands may be overlapped with kernel execution.

4.8.5.3.1.2 mem_advise

```
class handler {
 1
 2
      . . .
 3
      public:
 4
       . . .
       void mem_advise(const void *addr, size_t num_bytes, int advice);
 5
 6 };
 7
 8 class queue {
 9
     . . .
10
    public:
11
      . . .
12
      event mem_advise(const void *addr, size_t num_bytes, int advice);
13 };
```

Parameters

- void* addr address of allocation
- size_t num_bytes number of bytes in the allocation
- int advice device-defined advice for the specified allocation. A value of 0 reverts the advice for addr to the default behavior.

Return Value Returns an event representing the operation.

4.8.5.4 General

```
4.8.5.4.1 malloc
```

```
9 const sycl::device& dev,
10 const sycl::context& ctxt,
11 usm::alloc kind);
```

Parameters

- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- **const sycl::device&** dev the SYCL device to allocate on (if applicable)
- const sycl::context& ctxt the SYCL context to which device belongs
- usm::alloc kind the type of allocation to perform
- **Return value** Returns a pointer to the newly allocated kind memory on the specified device on success. If kind is alloc::host, dev is ignored. Memory allocated by sycl::malloc must be deallocated with sycl::free to avoid memory leaks. On failure, returns nullptr.

```
1
  (1)
2
   void *sycl::malloc(size_t num_bytes,
3
                      const sycl::queue& q,
4
                      usm::alloc kind);
5
  (2)
6
  template <typename T>
7
  T *sycl::malloc(size_t count,
8
                  const sycl::queue& q,
9
                   usm::alloc kind);
```

Parameters

- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::queue& q the SYCL q that provides the device (if applicable) and context to allocate against
- usm::alloc kind the type of allocation to perform

Return value Returns a pointer to the newly allocated kind memory on success. Memory allocated by sycl:: malloc must be deallocated with sycl::free to avoid memory leaks. On failure, returns nullptr.

4.8.5.4.2 aligned_alloc

```
1 (1)
2 void *sycl::aligned_alloc(size_t alignment,
3 size_t num_bytes,
4 const sycl::device& dev,
5 const sycl::context& ctxt,
6 usm::alloc kind);
7 (2)
```

Parameters

- size_t alignment specifies the byte alignment. Must be a valid alignment supported by the implementation.
- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- **const sycl::device&** dev the SYCL device to allocate on (if applicable)
- const sycl::context& ctxt the SYCL context to which device belongs
- usm::alloc kind the type of allocation to perform
- **Return value** Returns a pointer to the newly allocated kind memory on the specified device on success. If kind is alloc::host, dev is ignored. Memory allocated by sycl::aligned_alloc must be deallocated with sycl::free to avoid memory leaks. On failure, returns nullptr.

```
1
   (1)
2
   void *sycl::aligned_alloc(size_t alignment,
3
                              size_t num_bytes,
4
                              const sycl::queue& q,
5
                              usm::alloc kind);
6 (2)
7
   template <typename T>
8 T* sycl::aligned_alloc(size_t alignment,
9
                           size_t count,
10
                           const sycl::queue& q,
11
                           usm::alloc kind);
```

Parameters

- size_t alignment specifies the byte alignment. Must be a valid alignment supported by the implementation.
- (1) size_t num_bytes number of bytes to allocate
- (2) size_t count number of elements of type T to allocate
- const sycl::queue& q the SYCL q that provides the device (if applicable) and context to allocate against.
- usm::alloc kind the type of allocation to perform
- **Return value** Returns a pointer to the newly allocated kind memory on success. Memory allocated by sycl:: aligned_alloc must be deallocated with sycl::free to avoid memory leaks. On failure, returns nullptr.

4.8.5.4.3 free

1 void sycl::free(void* ptr, sycl::context& context);

Parameters

- void* ptr pointer to the memory to deallocate. Must have been allocated by a SYCL malloc or aligned_alloc function.
- const sycl::context& ctxt the SYCL context in which ptr was allocated

Return value none

1 void sycl::free(void* ptr, sycl::queue& q);

Parameters

- void* ptr pointer to the memory to deallocate. Must have been allocated by a SYCL malloc or aligned_alloc function.
- const sycl::queue& q the SYCL queue that provides the context in which ptr was allocated

Return value none

4.8.6 Unified shared memory information

4.8.6.1 Pointer queries

- 4.8.6.1.1 get_pointer_type
- 1 usm::alloc get_pointer_type(const void *ptr, const context &ctxt);

Parameters

- const void* ptr the pointer to query.
- const sycl::context& ctxt the SYCL context to which the USM allocation belongs
- **Return value** Returns the USM allocation type for ptr if ptr falls inside a valid USM allocation. If ctxt is a host context, returns usm::alloc::host. Returns usm::alloc::unknown if ptr is not a valid USM allocation.

4.8.6.1.2 get_pointer_device

1 sycl::device get_pointer_device(const void *ptr, const context &ctxt);

Parameters

- const void* ptr the pointer to query
- const sycl::context& ctxt the SYCL context to which the USM allocation belongs

Return value Returns the device associated with the USM allocation. If ctxt is a host context, returns the host device in ctxt. If ptr is an allocation of type usm::alloc::host, returns the first device in ctxt. Throws an error if ptr is not a valid USM allocation.

4.9 SYCL scheduling

SYCL 1.2.1 defines an execution model based on tasks submitted to Out-of-Order queues. Dependences between these tasks are constructed from the data they read and write. The data usage of a task is conveyed to the runtime by constructing accessors on buffer objects that specify their intent. Pointers obtained from using explicit memory management interfaces in SYCL cannot create accessors, so dependence graphs cannot be constructed in the same fashion. New member functions are required to specify dependences between tasks.

4.9.1 DAGs without accessors

Unified Shared Memory changes how the SYCL runtime manages data movement. Since the runtime might no longer be responsible for orchestrating data movement, it makes sense to enable a way to build dependence graphs based on ordering computations rather than accesses to data inside them. Conveniently, a SYCL queue already returns an event upon calls to submit. These events can be used by the programmer to wait for the submitted task to complete.

```
1 queue q;
2 auto dev = q.get_device();
 3
   auto ctxt = q.get_context();
   float* a = static_cast<float*>(malloc_shared(10*sizeof(float), dev, ctxt));
 4
 5
    float* b = static_cast<float*>(malloc_shared(10*sizeof(float), dev, ctxt));
 6
   float* c = static_cast<float*>(malloc_shared(10*sizeof(float), dev, ctxt));
 7
8
    auto e = q.submit([&](handler& cgh) {
9
      cgh.parallel_for(range<1> {10}, [=](id<1> ID) {
10
        size_t i = ID[0];
11
        c[i] = a[i] + b[i];
12
      });
13 });
14 e.wait();
```

4.9.2 Coarse grain DAGs with depends_on

While SYCL already defines the capability to wait on specific tasks, programmers should still be able to easily define relationships between tasks.

```
1 class handler {
2 ...
3 public:
4 ...
5 void depends_on(event e);
6 void depends_on(const std::vector<event> &e);
7 };
```

Parameters e - event or vector of events representing task(s) required to complete before this task may begin

Return value none

4.10 Expressing parallelism through kernels

4.10.1 Ranges and index space identifiers

The data parallelism of the SYCL kernel execution model requires instantiation of a parallel execution over a range of iteration space coordinates. To achieve this, SYCL exposes types to define the range of execution and to identify a given execution instance's point in the iteration space.

The following types are defined: range, nd_range, id, item, h_item, nd_item and group.

When constructing ids or ranges from integers, the elements are written in row-major format.

| Туре | Description |
|----------|---|
| id | A point within a range |
| range | Bounds over which an id may vary |
| item | Pairing of an id (specific point) and the |
| | range that it is bounded by |
| nd_range | Encapsules both global and local (work- |
| | group size) ranges over which work-item |
| | ids will vary |
| nd_item | Encapsulates two items, one for global id |
| | and range, and one for local id and range |
| h_item | Index point queries within hierarchical par- |
| | allelism (parallel_for_work_item). En- |
| | capsulates physical global and local ids and |
| | ranges, as well as a logical local id and |
| | range defined by hierarchical parallelism |
| group | Work-group queries within hierarchical par- |
| | allelism (parallel_for_work_group), and |
| | exposes the parallel_for_work_item con- |
| | struct that identifies code to be executed by |
| | each work-item. Encapsulates work-group |
| | ids and ranges |
| | End of table |

Table 4.73: Summary of types used to identify points in an index space, and ranges over which those points can vary.

4.10.1.1 range class

range<int dimensions> is a 1D, 2D or 3D vector that defines the iteration domain of either a single work-group in a parallel dispatch, or the overall dimensions of the dispatch. It can be constructed from integers.

The SYCL range class template provides the common by-value semantics (see Section 4.5.4).

A synopsis of the SYCL range class is provided below. The constructors, member functions and non-member functions of the SYCL range class are listed in Tables 4.74, 4.75 and 4.76 respectively. The additional common special member functions and common member functions are listed in 4.5.4 in Tables 4.3 and 4.4 respectively.

```
1 namespace sycl {
 2 template <int dimensions = 1>
 3 class range {
 4 public:
 5
      /* The following constructor is only available in the range class specialization where:
          dimensions==1 */
 6
      range(size_t dim0);
 7
      /* The following constructor is only available in the range class specialization where:
          dimensions==2 */
 8
      range(size_t dim0, size_t dim1);
 9
      /* The following constructor is only available in the range class specialization where:
          dimensions==3 */
10
      range(size_t dim0, size_t dim1, size_t dim2);
11
12
       /* -- common interface members -- */
13
14
      size_t get(int dimension) const;
15
      size_t &operator[](int dimension);
16
      size_t operator[](int dimension) const;
17
18
      size_t size() const;
19
20
      // OP is: +, -, *, /, %, <<, >>, &, |, ^, &&, ||, <, >, <=, >=
21
      friend range operatorOP(const range &lhs, const range &rhs) { /* ... */ }
22
      friend range operatorOP(const range &lhs, const size_t &rhs) { /* ... */ }
23
      // OP is: +=, -=, *=, /=, %=, <<=, >>=, &=, |=, ^=
24
25
      friend range & operatorOP(range &lhs, const range &rhs) { /* ... */ }
26
      friend range & operatorOP(range &lhs, const size_t &rhs) { /* ... */ }
27
28
      // OP is: +, -, *, /, %, <<, >>, &, |, ^, &&, ||, <, >, <=, >=
29
      friend range operatorOP(const size_t &lhs, const range &rhs) { /* ... */ }
30
31 };
32
33 // Deduction guides
34 range(size_t) -> range<1>;
35 range(size_t, size_t) -> range<2>;
36 range(size_t, size_t, size_t) -> range<3>;
37
38 } // sycl
```

| Constructor | Description |
|---|---|
| <pre>range(size_t dim0)</pre> | Construct a 1D range with value dim0. |
| | Only valid when the template parameter |
| | dimensions is equal to 1. |
| <pre>range(size_t dim0, size_t dim1)</pre> | Construct a 2D range with values dim0 and |
| | dim1. Only valid when the template param- |
| | eter dimensions is equal to 2. |
| <pre>range(size_t dim0, size_t dim1, size_t dim2)</pre> | Construct a 3D range with values dim0, |
| | dim1 and dim2. Only valid when the tem- |
| | plate parameter dimensions is equal to 3. |
| | End of table |

Table 4.74: Constructors of the range class template.

| Member function | Description | | |
|--|---|--|--|
| <pre>size_t get(int dimension)const</pre> | Return the value of the specified dimension | | |
| | of the range. | | |
| <pre>size_t &operator[](int dimension)</pre> | Return the l-value of the specified dimen- | | |
| | sion of the range. | | |
| <pre>size_t operator[](int dimension)const</pre> | Return the value of the specified dimension | | |
| | of the range. | | |
| <pre>size_t size()const</pre> | Return the size of the range computed as | | |
| | dimension0**dimensionN. | | |
| | End of table | | |

| Table 4.75: | Member | functions | of the | range | class | template. |
|-------------|--------|-----------|--------|-------|-------|-----------|
| | | | | | | ····· |

| Hidden friend function | Description |
|---|--|
| <pre>range operatorOP(const range &lhs, const range &rhs)</pre> | Where OP is: +, -, *, /, %, <<, >>, &, , ^, &&, |
| | , <, >, <=, >=. |
| | Constructs and returns a new instance of the |
| | SYCL range class template with the same |
| | dimensionality as 1hs range, where each el- |
| | ement of the new SYCL range instance is |
| | the result of an element-wise OP operator be- |
| | tween each element of 1hs range and each |
| | element of the rhs range. If the operator re- |
| | turns a bool the result is the cast to size_t. |
| <pre>range operatorOP(const range &lhs, const size_t &rhs</pre> | Where OP is: +, -, *, /, %, <<, >>, &, , ^, &&, |
|) | , <, >, <=, >=. |
| | Constructs and returns a new instance of the |
| | SYCL range class template with the same |
| | dimensionality as 1hs range, where each el- |
| | ement of the new SYCL range instance is |
| | the result of an element-wise OP operator be- |
| | tween each element of this SYCL range and |
| | the rhs size_t. If the operator returns a |
| | bool the result is the cast to size_t. |
| <pre>range &operatorOP(range &lhs, const range &rhs)</pre> | Where OP is: +=, -=,*=, /=, %=, <<=, >>=, &=, |
| | =, ^=. |
| | Assigns each element of 1hs range instance |
| | with the result of an element-wise OP opera- |
| | tor between each element of 1hs range and |
| | each element of the rhs range and returns |
| | lhs range. If the operator returns a bool the |
| | result is the cast to size_t. |
| | Continued on next page |

Table 4.76: Hidden friend functions of the SYCL range class template.

| Hidden friend function | Description |
|---|---|
| <pre>range &operatorOP(range &lhs, const size_t &rhs)</pre> | Where OP is: +=, -=, *=, /=, %=, <<=, >>=, &=, |
| | =, ^=. |
| | Assigns each element of 1hs range instance |
| | with the result of an element-wise OP opera- |
| | tor between each element of 1hs range and |
| | the rhs size_t and returns lhs range. If the |
| | operator returns a bool the result is the cast |
| | to size_t. |
| <pre>range operatorOP(const size_t &lhs, const range &rhs</pre> | Where OP is: +, -, *, /, %, <<, >>, &, , ^, &&, |
|) | , <, >, <=, >=. |
| | Constructs and returns a new instance of |
| | the SYCL range class template with the |
| | same dimensionality as the rhs SYCL range |
| | , where each element of the new SYCL |
| | range instance is the result of an element- |
| | wise OP operator between the lhs size_t |
| | and each element of the rhs SYCL range. If |
| | the operator returns a bool the result is the |
| | cast to size_t. |
| | End of table |

Table 4.76: Hidden friend functions of the SYCL range class template.

4.10.1.2 nd_range class

```
1 namespace sycl {
 2
   template <int dimensions = 1>
 3
   class nd_range {
 4
    public:
 5
 6
       /* -- common interface members -- */
 7
      nd_range(range<dimensions> globalSize, range<dimensions> localSize,
 8
 9
               id<dimensions> offset = id<dimensions>());
10
11
      range<dimensions> get_global_range() const;
12
      range<dimensions> get_local_range() const;
13
      range<dimensions> get_group_range() const;
14
      id<dimensions> get_offset() const;
15 };
16 } // namespace sycl
```

nd_range<int dimensions> defines the iteration domain of both the work-groups and the overall dispatch. To define this the nd_range comprises two ranges: the whole range over which the kernel is to be executed, and the range of each work group.

The SYCL nd_range class template provides the common by-value semantics (see Section 4.5.4).

A synopsis of the SYCL nd_range class is provided below. The constructors and member functions of the SYCL nd_range class are listed in Tables 4.77 and 4.78 respectively. The additional common special member functions
and common member functions are listed in 4.5.4 in Tables 4.3 and 4.4 respectively.

| Constructor | Description |
|--|---|
| <pre>nd_range<dimensions>(</dimensions></pre> | Construct an nd_range from the local and |
| <pre>range<dimensions> globalSize,</dimensions></pre> | global constituent ranges as well as an op- |
| <pre>range<dimensions> localSize)</dimensions></pre> | tional offset. If the offset is not provided it |
| <pre>id<dimensions> offset = id<dimensions>())</dimensions></dimensions></pre> | will default to no offset. |
| | End of table |

Table 4.77: Constructors of the nd_range class.

| Member function | Description |
|---|--|
| <pre>range<dimensions> get_global_range()const</dimensions></pre> | Return the constituent global range. |
| <pre>range<dimensions> get_local_range()const</dimensions></pre> | Return the constituent local range. |
| <pre>range<dimensions> get_group_range()const</dimensions></pre> | Return a range representing the number of |
| | groups in each dimension. This range would |
| | result from globalSize/localSize as pro- |
| | vided on construction. |
| <pre>id<dimensions> get_offset()const</dimensions></pre> | Return the constituent offset. |
| | End of table |

Table 4.78: Member functions for the nd_range class.

4.10.1.3 id class

id<int dimensions> is a vector of dimensions that is used to represent an id into a global or local range. It can be used as an index in an accessor of the same rank. The [n] operator returns the component n as an size_t.

The SYCL id class template provides the common by-value semantics (see Section 4.5.4).

A synopsis of the SYCL id class is provided below. The constructors, member functions and non-member functions of the SYCL id class are listed in Tables 4.79, 4.80 and 4.81 respectively. The additional common special member functions and common member functions are listed in 4.5.4 in Tables 4.3 and 4.4 respectively.

```
1 namespace sycl {
2
   template <int dimensions = 1>
3
   class id {
   public:
4
5
     id();
6
     /* The following constructor is only available in the id class
7
      * specialization where: dimensions==1 */
8
9
     id(size_t dim0);
      /* The following constructor is only available in the id class
10
      * specialization where: dimensions==2 */
11
12
      id(size_t dim0, size_t dim1);
      /* The following constructor is only available in the id class
13
14
      * specialization where: dimensions==3 */
15
     id(size_t dim0, size_t dim1, size_t dim2);
16
       /* -- common interface members -- */
17
```

```
18
19
      id(const range<dimensions> &range);
      id(const item<dimensions> &item);
20
21
22
      size_t get(int dimension) const;
23
      size_t &operator[](int dimension);
24
      size_t operator[](int dimension) const;
25
26
      // OP is: +, -, *, /, %, <<, >>, &, |, ^, &&, ||, <, >, <=, >=
27
        friend id operatorOP(const id &lhs, const id &rhs) { /* ... */ }
28
        friend id operatorOP(const id &lhs, const size_t &rhs) { /* ... */ }
29
      // OP is: +=, -=, *=, /=, %=, <<=, >>=, &=, |=, ^=
30
        friend id &operatorOP(id &lhs, const id &rhs) { /* ... */ }
31
32
        friend id &operatorOP(id &lhs, const size_t &rhs) { /* ... */ }
33
34
      // OP is: +, -, *, /, %, <<, >>, &, |, ^, &&, ||, <, >, <=, >=
35
        friend id operatorOP(const size_t &lhs, const id &rhs) { /* ... */ }
36
37 };
38
39 // Deduction guides
40 id(size_t) -> id<1>;
41 id(size_t, size_t) -> id<2>;
42 id(size_t, size_t, size_t) -> id<3>;
43
44 } // namespace sycl
```

| Constructor | Description |
|--|---|
| id() | Construct a SYCL id with the value 0 for |
| | each dimension. |
| <pre>id(size_t dim0)</pre> | Construct a 1D id with value dim0. |
| | Only valid when the template parameter |
| | dimensions is equal to 1. |
| <pre>id(size_t dim0, size_t dim1)</pre> | Construct a 2D id with values dim0, dim1. |
| | Only valid when the template parameter |
| | dimensions is equal to 2. |
| <pre>id(size_t dim0, size_t dim1, size_t dim2)</pre> | Construct a 3D id with values dim0, dim1, |
| | dim2. Only valid when the template param- |
| | eter dimensions is equal to 3. |
| <pre>id(const range<dimensions> ⦥)</dimensions></pre> | Construct an id from the dimensions of |
| | range. |
| <pre>id(const item<dimensions> &item)</dimensions></pre> | Construct an id from item.get_id(). |
| | End of table |

Table 4.79: Constructors of the *id* class template.

| Member function | Description |
|--|---|
| <pre>size_t get(int dimension)const</pre> | Return the value of the id for dimension |
| | dimension. |
| <pre>size_t &operator[](int dimension)</pre> | Return a reference to the requested dimen- |
| | sion of the id object. |
| <pre>size_t operator[](int dimension)const</pre> | Return the value of the requested dimension |
| | of the id object. |
| | End of table |

Table 4.80: Member functions of the *id* class template.

| Hidden friend function | Description |
|--|--|
| <pre>id operatorOP(const id &lhs, const id &rhs)</pre> | Where OP is: +, -, *, /, %, <<, >>, &, , ^, &&, |
| | , <, >, <=, >=. |
| | Constructs and returns a new instance of the |
| | SYCL id class template with the same di- |
| | mensionality as 1hs id, where each element |
| | of the new SYCL id instance is the result of |
| | an element-wise OP operator between each |
| | element of 1hs id and each element of the |
| | rhs id. If the operator returns a bool, the |
| | result is the cast to size_t. |
| <pre>id operatorOP(const id &lhs, const size_t &rhs)</pre> | Where OP is: +, -, *, /, %, <<, >>, &, , ^, &&, |
| | , <, >, <=, >=. |
| | Constructs and returns a new instance of the |
| | SYCL id class template with the same di- |
| | mensionality as 1hs id, where each element |
| | of the new SYCL id instance is the result of |
| | an element-wise OP operator between each |
| | element of lhs id and the rhs size_t. If the |
| | operator returns a bool, the result is the cast |
| | to size_t. |
| id &operatorOP(id &lhs, const id &rhs) | Where OP 1s: +=, -=,*=, /=, %=, <<=, >>=, &=, |
| | |
| | Assigns each element of this id instance |
| | with the result of an element-wise op opera- |
| | tor between each element of Ins 1d and each |
| | the operator returns a heal, the result is the |
| | cost to cize t |
| id forematerOP(id flbg, const size t frbg) | $\frac{\text{Cast to Size_t.}}{\text{Where OP is: } * - / - \% - / - \% - }$ |
| iu aoperatoror(iu ains, const size_t ains) | W HELE OF IS. +-,, "-, /-, /o-, <<-, >>-, d-, |
| | Assigns each element of the id instance |
| | with the result of an element-wise OP opera- |
| | tor between each element of the id and the |
| | rhs size t and returns lhs id. If the oper- |
| | ator returns a bool, the result is the cast to |
| | size_t. |
| | Continued on next page |

Table 4.81: Hidden friend functions of the *id* class template.

| Hidden friend function | Description |
|--|--|
| <pre>id operatorOP(const size_t &lhs, const id &rhs)</pre> | Where OP is: +, -, *, /, %, <<, >>, &, , ^, &&, |
| | , <, >, <=, >=. |
| | Constructs and returns a new instance of the |
| | SYCL id class template with the same di- |
| | mensionality as the rhs SYCL id, where |
| | each element of the new SYCL id instance |
| | is the result of an element-wise OP operator |
| | between the lhs size_t and each element of |
| | the rhs SYCL id. If the operator returns a |
| | bool, the result is the cast to size_t. |
| | End of table |

Table 4.81: Hidden friend functions of the id class template.

4.10.1.4 item class

item identifies an instance of the function object executing at each point in a range. It is passed to a parallel_for call or returned by member functions of h_item. It encapsulates enough information to identify the work-item's range of possible values and its ID in that range. It can optionally carry the offset of the range if provided to the parallel_for. Instances of the item class are not user-constructible and are passed by the runtime to each instance of the function object.

The SYCL item class template provides the common by-value semantics (see Section 4.5.4).

A synopsis of the SYCL item class is provided below. The member functions of the SYCL item class are listed in Table 4.80. The additional common special member functions and common member functions are listed in 4.5.4 in Tables 4.3 and 4.4 respectively.

```
1 namespace sycl {
 2 template <int dimensions = 1, bool with_offset = true>
 3
    class item {
 4
    public:
 5
      item() = delete;
 6
 7
       /* -- common interface members -- */
 8
 9
      id<dimensions> get_id() const;
10
11
      size_t get_id(int dimension) const;
12
13
      size_t operator[](int dimension) const;
14
15
      range<dimensions> get_range() const;
16
17
      size_t get_range(int dimension) const;
18
19
      // only available if with_offset is true
20
      id<dimensions> get_offset() const;
21
22
      // only available if with_offset is false
23
      operator item<dimensions, true>() const;
```

24 25 // only available if dimensions == 1 26 operator size_t() const; 27 28 size_t get_linear_id() const; 29 };

```
30 } // namespace sycl
```

| Member function | Description |
|--|---|
| <pre>id<dimensions> get_id()const</dimensions></pre> | Return the constituent id representing the |
| | work-item's position in the iteration space. |
| <pre>size_t get_id(int dimension)const</pre> | Return the same value as get_id()[|
| | dimension]. |
| <pre>size_t operator[](int dimension)const</pre> | Return the same value as get_id(|
| | dimension). |
| <pre>range<dimensions> get_range()const</dimensions></pre> | Returns a range representing the dimen- |
| | sions of the range of possible values of the |
| | item. |
| <pre>size_t get_range(int dimension)const</pre> | Return the same value as get_range().get |
| | (dimension). |
| <pre>id<dimensions> get_offset()const</dimensions></pre> | Returns an id representing the <i>n</i> - |
| | dimensional offset provided to the |
| | parallel_for and that is added by the |
| | runtime to the global-ID of each work-item, |
| | if this item represents a global range. For an |
| | item converted from an item with no offset |
| | this will always return an id of all 0 values. |
| | This member function is only available if |
| | with_offset is true. |
| <pre>operator item<dimensions, true="">()const</dimensions,></pre> | Available only when: with_offset == |
| | false |
| | Returns an item representing the same in- |
| | formation as the object holds but also in- |
| | cludes the offset set to 0. This conversion |
| | allow users to seamlessly write code that as- |
| | sumes an onset and sum provides an onset- |
| | Available only when dimensions 1 |
| operator size_t()const | Available only when unerstons == 1 Deturns the index representing the work |
| | item position in the iteration space |
| | nem position in the neration space. |
| <pre>size_t get_linear_id()const</pre> | Return the id as a linear index value. Cal- |
| | culating a linear address from the multi- |
| | dimensional index follow the equation 4.3. |
| | End of table |

Table 4.82: Member functions for the item class.

4.10.1.5 nd_item class

nd_item<int dimensions> identifies an instance of the function object executing at each point in an nd_range< int dimensions> passed to a parallel_for call. It encapsulates enough information to identify the work-item's local and global ids, the work-group id and also provides access to the group and sub_group classes. Instances of the nd_item<int dimensions> class are not user-constructible and are passed by the runtime to each instance of the function object.

The SYCL nd_item class template provides the common by-value semantics (see Section 4.5.4).

A synopsis of the SYCL nd_item class is provided below. The member functions of the SYCL nd_item class are listed in Table 4.83. The additional common special member functions and common member functions are listed in 4.5.4 in Tables 4.3 and 4.4 respectively.

```
1 namespace sycl {
 2 template <int dimensions = 1>
 3
    class nd_item {
 4
    public:
 5
      nd_item() = delete;
 6
 7
       /* -- common interface members -- */
 8
 9
      id<dimensions> get_global_id() const;
10
11
      size_t get_global_id(int dimension) const;
12
13
      size_t get_global_linear_id() const;
14
15
      id<dimensions> get_local_id() const;
16
17
      size_t get_local_id(int dimension) const;
18
19
      size_t get_local_linear_id() const;
20
21
      group<dimensions> get_group() const;
22
23
      size_t get_group(int dimension) const;
24
25
      size_t get_group_linear_id() const;
26
27
      range<dimensions> get_group_range() const;
28
29
      size_t get_group_range(int dimension) const;
30
31
      range<dimensions> get_global_range() const;
32
33
      size_t get_global_range(int dimension) const;
34
35
      range<dimensions> get_local_range() const;
36
37
      size_t get_local_range(int dimension) const;
38
39
      id<dimensions> get_offset() const;
40
41
      nd_range<dimensions> get_nd_range() const;
```

```
42
43
      template <typename dataT>
44
      device_event async_work_group_copy(decorated_local_ptr<dataT> dest,
45
        decorated_global_ptr<dataT> src, size_t numElements) const;
46
47
      template <typename dataT>
48
      device_event async_work_group_copy(decorated_global_ptr<dataT> dest,
49
        decorated_local_ptr<dataT> src, size_t numElements) const;
50
51
      template <typename dataT>
52
      device_event async_work_group_copy(decorated_local_ptr<dataT> dest,
53
        decorated_global_ptr<dataT> src, size_t numElements, size_t srcStride) const;
54
55
      template <typename dataT>
56
      device_event async_work_group_copy(decorated_global_ptr<dataT> dest,
57
        decorated_local_ptr<dataT> src, size_t numElements, size_t destStride) const;
58
59
      template <typename... eventTN>
60
      void wait_for(eventTN... events) const;
61 };
62 } // namespace sycl
```

| Member function | Description |
|---|--|
| <pre>id<dimensions> get_global_id()const</dimensions></pre> | Return the constituent global id represent- |
| | ing the work-item's position in the global it- |
| | eration space. |
| <pre>size_t get_global_id(int dimension)const</pre> | Return the constituent element of the global |
| | id representing the work-item's position in |
| | the nd-range in the given dimension. |
| <pre>size_t get_global_linear_id()const</pre> | Return the flattened id of the current work- |
| | item after subtracting the offset. Calculating |
| | a linear id from a multi-dimensional index |
| | follows the equation 4.3. |
| <pre>id<dimensions> get_local_id()const</dimensions></pre> | Return the constituent local id representing |
| | the work-item's position within the current |
| | work-group. |
| <pre>size_t get_local_id(int dimension)const</pre> | Return the constituent element of the lo- |
| | cal id representing the work-item's position |
| | within the current work-group in the given |
| | dimension. |
| <pre>size_t get_local_linear_id()const</pre> | Return the flattened id of the current work- |
| | item within the current work-group. Cal- |
| | culating a linear address from a multi- |
| | dimensional index follows the equation 4.3. |
| <pre>group<dimensions> get_group()const</dimensions></pre> | Return the constituent work-group, group |
| | representing the work-group's position |
| | within the overall nd-range. |
| <pre>sub_group get_sub_group()const</pre> | Return a sub_group representing the sub- |
| | group to which the work-item belongs. |
| | Continued on next page |

Table 4.83: Member functions for the nd_item class.

| Member function | Description |
|---|--|
| <pre>size_t get_group(int dimension)const</pre> | Return the constituent element of the group |
| | id representing the work-group's position |
| | within the overall nd_range in the given |
| | dimension. |
| <pre>size_t get_group_linear_id()const</pre> | Return the group id as a linear index value. |
| | Calculating a linear address from a multi- |
| | dimensional index follows the equation 4.3. |
| <pre>range<dimensions> get_group_range()const</dimensions></pre> | Returns the number of work-groups in the |
| | iteration space. |
| <pre>size_t get_group_range(int dimension)const</pre> | Return the number of work-groups for |
| | dimension in the iteration space. |
| <pre>range<dimensions> get_global_range()const</dimensions></pre> | Returns a range representing the dimen- |
| | sions of the global iteration space. |
| <pre>size_t get_global_range(int dimension)const</pre> | Return the same value as get_global |
| | range().get(dimension) |
| <pre>range<dimensions> get_local_range()const</dimensions></pre> | Returns a range representing the dimen- |
| | sions of the current work-group. |
| <pre>size_t get_local_range(int dimension)const</pre> | Return the same value as get_local |
| | range().get(dimension) |
| <pre>id<dimensions> get_offset()const</dimensions></pre> | Returns an id representing the n- |
| | dimensional offset provided to the con- |
| | structor of the nd_range and that is added |
| | by the runtime to the global id of each |
| | work-item. |
| <pre>nd_range<dimensions> get_nd_range()const</dimensions></pre> | Returns the nd_range of the current execu- |
| | tion. |
| <pre>template <typename datat=""></typename></pre> | Permitted types for dataT are all scalar and |
| <pre>device_event async_work_group_copy(</pre> | vector types. Asynchronously copies a num- |
| <pre>decorated_local_ptr<datat> dest,</datat></pre> | ber of elements specified by numElements |
| <pre>decorated_global_ptr<datat> src,</datat></pre> | from the source pointer src to desti- |
| <pre>size_t numElements)const</pre> | nation pointer dest and returns a SYCL |
| | device_event which can be used to wait on |
| | the completion of the copy. |
| <pre>template <typename datat=""></typename></pre> | Permitted types for dataT are all scalar and |
| <pre>device_event async_work_group_copy(</pre> | vector types. Asynchronously copies a num- |
| <pre>decorated_global_ptr<datat> dest,</datat></pre> | ber of elements specified by numElements |
| <pre>decorated_local_ptr<datat> src,</datat></pre> | trom the source pointer src to desti- |
| <pre>size_t numElements)const</pre> | nation pointer dest and returns a SYCL |
| | device_event which can be used to wait on |
| | the completion of the copy. |
| | Continued on next page |

Table 4.83: Member functions for the nd_item class.

| Member function | Description |
|--|--|
| <pre>template <typename datat=""></typename></pre> | Permitted types for dataT are all scalar and |
| <pre>device_event async_work_group_copy(</pre> | vector types. Asynchronously copies a num- |
| <pre>decorated_local_ptr<datat> dest,</datat></pre> | ber of elements specified by numElements |
| <pre>decorated_global_ptr<datat> src,</datat></pre> | from the source pointer src to destina- |
| <pre>size_t numElements, size_t srcStride)const</pre> | tion pointer dest with a source stride spec- |
| | ified by srcStride and returns a SYCL |
| | device_event which can be used to wait on |
| | the completion of the copy. |
| <pre>template <typename datat=""></typename></pre> | Permitted types for dataT are all scalar and |
| <pre>device_event async_work_group_copy(</pre> | vector types. Asynchronously copies a num- |
| <pre>decorated_global_ptr<datat> dest,</datat></pre> | ber of elements specified by numElements |
| <pre>decorated_local_ptr<datat> src,</datat></pre> | from the source pointer src to destination |
| <pre>size_t numElements, size_t destStride)const</pre> | pointer dest with a destination stride spec- |
| | ified by destStride and returns a SYCL |
| | device_event which can be used to wait on |
| | the completion of the copy. |
| <pre>template <typename eventtn=""></typename></pre> | Permitted type for eventTN is device_event. |
| <pre>void wait_for(eventTN events)const</pre> | Waits for the asynchronous operations asso- |
| | ciated with each device_event to complete. |
| | End of table |

Table 4.83: Member functions for the nd_item class.

4.10.1.6 h_item class

h_item<int dimensions> identifies an instance of a group::parallel_for_work_item function object executing at each point in a local range<int dimensions> passed to a parallel_for_work_item call or to the corresponding parallel_for_work_group call if no range is passed to the parallel_for_work_item call. It encapsulates enough information to identify the work-item's local and global items according to the information given to parallel_for_work_group (physical ids) as well as the work-item's logical local items in the logical local range. All returned items objects are offset-less. Instances of the h_item<int dimensions> class are not user-constructible and are passed by the runtime to each instance of the function object.

The SYCL h_item class template provides the common by-value semantics (see Section 4.5.4).

A synopsis of the SYCL h_item class is provided below. The member functions of the SYCL h_item class are listed in Table 4.84. The additional common special member functions and common member functions are listed in 4.5.4 in Tables 4.3 and 4.4 respectively.

```
1 namespace sycl {
2 template <int dimensions>
3 class h_item {
4
   public:
5
     h_item() = delete;
6
7
      /* -- common interface members -- */
8
9
      item<dimensions, false> get_global() const;
10
11
      item<dimensions, false> get_local() const;
```

| 12 | |
|-------------|---|
| 13 | <pre>item<dimensions, false=""> get_logical_local() const;</dimensions,></pre> |
| 14 | |
| 15 | <pre>item<dimensions, false=""> get_physical_local() const;</dimensions,></pre> |
| 16 | |
| 17 | <pre>range<dimensions> get_global_range() const;</dimensions></pre> |
| 18 | |
| 19 | <pre>size_t get_global_range(int dimension) const;</pre> |
| 20 | |
| 21 | <pre>id<dimensions> get_global_id() const;</dimensions></pre> |
| 22 | |
| 23 | <pre>size_t get_global_id(int dimension) const;</pre> |
| 24 | |
| 25 | <pre>range<dimensions> get_local_range() const;</dimensions></pre> |
| 26 | |
| 27 | <pre>size_t get_local_range(int dimension) const;</pre> |
| 28 | |
| 29 | <pre>id<dimensions> get_local_id() const;</dimensions></pre> |
| 30 | |
| 31 | <pre>size_t get_local_id(int dimension) const;</pre> |
| 32 | |
| 33 | <pre>range<dimensions> get_logical_local_range() const;</dimensions></pre> |
| 34 | |
| 35 | <pre>size_t get_logical_local_range(int dimension) const;</pre> |
| 36 | |
| 37 | <pre>id<dimensions> get_logical_local_id() const;</dimensions></pre> |
| 38 | |
| 39 | <pre>size_t get_logical_local_id(int dimension) const;</pre> |
| 40 | and dimensions, set abusical local analy () south |
| 41 | range <dimensions> get_physical_local_range() const;</dimensions> |
| 42 | size t est abusies] less and (int dimension) esset. |
| 43 | size_t get_physical_local_range(int dimension) const; |
| 44 | id dimensions, and physical local id () constru |
| 45 | id dimensions get_physical_id() const, |
| 40 | size t get physical local id(int dimension) const: |
| -+ / / 8 | |
| 40 40 | }. |
| 50 | <pre>> // namesnace svcl</pre> |
| 50 | J // nuncspace Syce |

| Member function | Description |
|---|---|
| <pre>item<dimensions, false=""> get_global()const</dimensions,></pre> | Return the constituent global item rep- resenting the work-item's position in the global iteration space as provided upon ker- nel invocation. |
| <pre>item<dimensions, false=""> get_local()const</dimensions,></pre> | Return the same value as get_logical_local(). |
| | Continued on next page |

Table 4.84: Member functions for the h_item class.

| Member function | Description |
|--|--|
| <pre>item<dimensions, false=""> get_logical_local()const</dimensions,></pre> | Return the constituent element of the |
| | logical local item work-item's position |
| | in the local iteration space as provided |
| | upon the invocation of the group:: |
| | parallel for work item. |
| | If the group::parallel for work item |
| | was called without any logical local range |
| | then the member function returns the |
| | physical local item. |
| | A physical id can be computed from a |
| | logical id by getting the remainder of the |
| | integer division of the logical id and the |
| | physical range: get logical local().get |
| | ()% get physical local.get range()== |
| | get physical local() get() |
| item <dimensions false=""> get physical local()const</dimensions> | Return the constituent element of the physi- |
| icem (dimensions), faibes get_physical_focul()const | cal local item work-item's position in the lo- |
| | cal iteration space as provided (by the user |
| | or the runtime) upon the kernel invocation. |
| range <dimensions> get global range()const</dimensions> | Return the same value as get global() |
| Tunge (universitions) get_grobut_runge() const | get range() |
| size t get global range(int dimension)const | Return the same value as get global() |
| Sinc_t get_grobar_range(int armenoron) const | get range(dimension) |
| id <dimensions> get global id()const</dimensions> | Return the same value as get global() |
| | get id() |
| size t get global id(int dimension)const | Return the same value as get global(). |
| | <pre>get_id(dimension)</pre> |
| <pre>range<dimensions> get_local_range()const</dimensions></pre> | Return the same value as get_local(). |
| | get_range() |
| <pre>size_t get_local_range(int dimension)const</pre> | Return the same value as get_local(). |
| | get_range(dimension) |
| <pre>id<dimensions> get_local_id()const</dimensions></pre> | Return the same value as get_local(). |
| | <pre>get_id()</pre> |
| <pre>size_t get_local_id(int dimension)const</pre> | Return the same value as get_local(). |
| | get_id(dimension) |
| <pre>range<dimensions> get_logical_local_range()const</dimensions></pre> | Return the same value as |
| | <pre>get_logical_local().get_range()</pre> |
| <pre>size_t get_logical_local_range(int dimension)const</pre> | Return the same value as |
| | <pre>get_logical_local().get_range(</pre> |
| | dimension) |
| <pre>id<dimensions> get_logical_local_id()const</dimensions></pre> | Return the same value as |
| | <pre>get_logical_local().get_id()</pre> |
| <pre>size_t get_logical_local_id(int dimension)const</pre> | Return the same value as |
| | <pre>get_logical_local().get_id(dimension)</pre> |
| <pre>range<dimensions> get_physical_local_range()const</dimensions></pre> | Return the same value as |
| | <pre>get_physical_local().get_range()</pre> |
| | Continued on next page |

Table 4.84: Member functions for the h_item class.

| Member function | Description |
|---|--|
| <pre>size_t get_physical_local_range(int dimension)const</pre> | Return the same value as |
| | <pre>get_physical_local().get_range(</pre> |
| | dimension) |
| <pre>id<dimensions> get_physical_local_id()const</dimensions></pre> | Return the same value as |
| | <pre>get_physical_local().get_id()</pre> |
| <pre>size_t get_physical_local_id(int dimension)const</pre> | Return the same value as |
| | <pre>get_physical_local().get_id(dimension</pre> |
| |) |
| | End of table |

Table 4.84: Member functions for the h_item class.

4.10.1.7 group class

The group<int dimensions> encapsulates all functionality required to represent a particular work-group within a parallel execution. It is not user-constructable.

The local range stored in the group class is provided either by the programmer, when it is passed as an optional parameter to parallel_for_work_group, or by the runtime system when it selects the optimal work-group size. This allows the developer to always know how many concurrent work-items are active in each executing work-group, even through the abstracted iteration range of the parallel_for_work_item loops.

The SYCL group class template provides the common by-value semantics (see Section 4.5.4).

A synopsis of the SYCL group class is provided below. The member functions of the SYCL group class are listed in Table 4.85. The additional common special member functions and common member functions are listed in 4.5.4 in Tables 4.3 and 4.4 respectively.

```
1 namespace sycl {
 2 template <int Dimensions = 1>
 3 class group {
 4 public:
 5
 6
      using id_type = id<Dimensions>;
 7
      using range_type = range<Dimensions>;
 8
      using linear_id_type = size_t;
 9
      static constexpr int dimensions = Dimensions;
10
      static constexpr memory_scope fence_scope = memory_scope::work_group;
11
       /* -- common interface members -- */
12
13
14
      id<Dimensions> get_group_id() const;
15
16
      size_t get_group_id(int dimension) const;
17
18
      id<Dimensions> get_local_id() const;
19
20
      size_t get_local_id(int dimension) const;
21
22
      range<Dimensions> get_local_range() const;
23
24
      size_t get_local_range(int dimension) const;
```

```
25
26
      range<Dimensions> get_group_range() const;
27
28
      size_t get_group_range(int dimension) const;
29
30
      range<Dimensions> get_max_local_range() const;
31
32
      range<Dimensions> get_uniform_group_range() const;
33
     size_t operator[](int dimension) const;
34
35
36
      size_t get_group_linear_id() const;
37
38
      size_t get_local_linear_id() const;
39
40
      size_t get_group_linear_range() const;
41
42
      size_t get_local_linear_range() const;
43
44
      bool leader() const;
45
46
      template<typename workItemFunctionT>
47
      void parallel_for_work_item(const workItemFunctionT &func) const;
48
49
      template<typename workItemFunctionT>
50
      void parallel_for_work_item(range<dimensions> logicalRange,
51
        const workItemFunctionT &func) const;
52
53
      template <typename dataT>
54
      device_event async_work_group_copy(decorated_local_ptr<dataT> dest,
55
        decorated_global_ptr<dataT> src, size_t numElements) const;
56
57
      template <typename dataT>
58
      device_event async_work_group_copy(decorated_global_ptr<dataT> dest,
59
        decorated_local_ptr<dataT> src, size_t numElements) const;
60
61
      template <typename dataT>
62
      device_event async_work_group_copy(decorated_local_ptr<dataT> dest,
63
        decorated_global_ptr<dataT> src, size_t numElements, size_t srcStride) const;
64
65
      template <typename dataT>
66
     device_event async_work_group_copy(decorated_global_ptr<dataT> dest,
67
        decorated_local_ptr<dataT> src, size_t numElements, size_t destStride) const;
68
69
     template <typename... eventTN>
70
     void wait_for(eventTN... events) const;
71 };
72 } // sycl
```

| Member function | Description |
|--|--|
| <pre>id<dimensions> get_group_id()const</dimensions></pre> | Return an id representing the index of the |
| | work-group within the nd-range for every |
| | dimension. |
| <pre>size_t get_group_id(int dimension)const</pre> | Return the same value as get_id()[|
| | dimension]. |
| <pre>id<dimensions> get_local_id()const</dimensions></pre> | Return a SYCL id representing the calling |
| | work-item's position within the work-group. |
| | It is undefined behavior for this mem- |
| | ber function to be invoked from within a |
| | <pre>parallel_for_work_item context.</pre> |
| <pre>size_t get_local_id(int dimension)const</pre> | Return the calling work-item's position |
| | within the work-group in the specified di- |
| | mension. |
| | It is undefined behavior for this mem- |
| | ber function to be invoked from within a |
| | <pre>parallel_for_work_item context.</pre> |
| <pre>range<dimensions> get_local_range()const</dimensions></pre> | Return a SYCL range representing all di- |
| | mensions of the local range. This local range |
| | may have been provided by the programmer, |
| | or chosen by the SYCL runtime. |
| <pre>size_t get_local_range(int dimension)const</pre> | Return the dimension of the local range |
| | specified by the dimension parameter. |
| <pre>range<dimensions> get_group_range()const</dimensions></pre> | Return a range representing the number of |
| | work-groups in the nd_range. |
| <pre>size_t get_group_range(int dimension)const</pre> | Return element dimension from the con- |
| | stituent group range. |
| <pre>size_t operator[](int dimension)const</pre> | Return the same value as get_id(|
| | dimension). |
| <pre>range<dimensions> get_max_local_range()const</dimensions></pre> | Return a range representing the maximum |
| | number of work-items in any work-group in |
| | the nd_range. |
| <pre>range<dimensions> get_uniform_group_range()const</dimensions></pre> | Return a range representing the number of |
| | work-groups in the uniform region of the |
| | nd_range. |
| <pre>size_t get_group_linear_id()const</pre> | Get a linearized version of the work-group |
| | id. Calculating a linear work-group id from |
| | a multi-dimensional index follows the equa- |
| | tion 4.3. |
| <pre>size_t get_group_linear_range()const</pre> | Return the total number of work-groups in |
| | the nd_range. |
| <pre>size_t get_local_linear_id()const</pre> | Get a linearized version of the calling work- |
| | item's local id. Calculating a linear local id |
| | from a multi-dimensional index follows the |
| | equation 4.3. |
| | It is undefined behavior for this mem- |
| | ber function to be invoked from within a |
| | parallel_for_work_item context. |
| | Continued on next page |

| Member function | Description |
|--|---|
| <pre>size_t get_local_linear_range()const</pre> | Return the total number of work-items in |
| | the work-group. |
| <pre>bool leader()const</pre> | Return true for exactly one work-item in the |
| | work-group, if the calling work-item is the |
| | leader of the work-group, and false for all |
| | other work-items in the work-group. |
| | The leader of the work-group is determined |
| | during construction of the work-group, and |
| | is invariant for the lifetime of the work- |
| | group. The leader of the work-group is guar- |
| | anteed to be the work-item with a local id of |
| | 0. |
| <pre>template <typename workitemfunctiont=""></typename></pre> | Launch the work-items for this work-group. |
| <pre>void parallel_for_work_item(const</pre> | func is a function object type with a |
| workItemFunctionT &func)const | public member function void F::operator |
| | ()(h_item <dimensions>) representing the</dimensions> |
| | work-item computation. |
| | This member function can only be invoked |
| | within a parallel_for_work_group context. |
| | It is undefined behavior for this member |
| | function to be invoked from within the |
| | <pre>parallel_for_work_group form that does</pre> |
| | not define work-group size, because then the |
| | number of work-items that should execute |
| | the code is not defined. It is expected that |
| | this form of parallel_for_work_item is in- |
| | voked within the parallel_for_work_group |
| | form that specifies the size of a work-group. |
| | Continued on next page |

| Member function | Description |
|---|---|
| <pre>template <typename workitemfunctiont=""></typename></pre> | Launch the work-items for this work-group |
| <pre>void parallel_for_work_item(range<dimensions></dimensions></pre> | using a logical local range. The function ob- |
| <pre>logicalRange, const workItemFunctionT &func)</pre> | ject func is executed as if the kernel were in- |
| const | voked with logicalRange as the local range. |
| | This new local range is emulated and may |
| | not map one-to-one with the physical range. |
| | logicalRange is the new local range to be |
| | used. This range can be smaller or larger |
| | than the one used to invoke the kernel. func |
| | is a function object type with a public mem- |
| | <pre>ber function void F::operator()(h_item<</pre> |
| | dimensions>) representing the work-item |
| | computation. |
| | Note that the logical range does not need to |
| | be uniform across all work-groups in a ker- |
| | nel. For example the logical range may de- |
| | pend on a work-group varying query (e.g. |
| | <pre>group::get_linear_id), such that different</pre> |
| | work-groups in the same kernel invocation |
| | execute different logical range sizes. |
| | This member function can only be invoked |
| | within a parallel_for_work_group context. |
| template <typename datat=""></typename> | Permitted types for dataT are all scalar and |
| device_event async_work_group_copy(| vector types. Asynchronously copies a num- |
| <pre>decorated_local_ptr<datat> dest,</datat></pre> | ber of elements specified by numElements |
| decorated_global_ptr <datal> src,</datal> | from the source pointer src to desti- |
| SIZE_t numElements)const | had not pointer dest and returns a SYCL |
| | the completion of the conv |
| townlate (twomen date) | Dermitted types for date T are all scalar and |
| device event some vork group conv(| vector types. A synchronously copies a num |
| decorated global ntr(dataT> dest | ber of elements specified by numElements |
| decorated local ntr/dataT> src | from the source pointer src to desti- |
| size t numElements) const | nation pointer dest and returns a SVCI |
| SIZE_t numErement(S)CONSt | device event which can be used to wait on |
| | the completion of the conv |
| template <typename datat=""></typename> | Permitted types for dataT are all scalar and |
| device event async work group conv(| vector types Asynchronously copies a num- |
| decorated local ptr <datat> dest.</datat> | ber of elements specified by numElements |
| <pre>decorated_global_ptr<datat> src.</datat></pre> | from the source pointer src to destina- |
| size_t numElements, size t srcStride)const | tion pointer dest with a source stride spec- |
| ,,, ,, ,, ,, ,, ,, ,, ,, ,, ,, ,, ,, | ified by srcStride and returns a SYCL |
| | device_event which can be used to wait on |
| | the completion of the copy. |
| | Continued on next page |

| Member function | Description |
|--|--|
| <pre>template <typename datat=""></typename></pre> | Permitted types for dataT are all scalar and |
| <pre>device_event async_work_group_copy(</pre> | vector types. Asynchronously copies a num- |
| <pre>decorated_global_ptr<datat> dest,</datat></pre> | ber of elements specified by numElements |
| <pre>decorated_local_ptr<datat> src,</datat></pre> | from the source pointer src to destination |
| <pre>size_t numElements, size_t destStride)const</pre> | pointer dest with a destination stride spec- |
| | ified by destStride and returns a SYCL |
| | device_event which can be used to wait on |
| | the completion of the copy. |
| <pre>template <typename eventtn=""></typename></pre> | Permitted type for eventTN is device_event. |
| <pre>void wait_for(eventTN events)const</pre> | Waits for the asynchronous operations asso- |
| | ciated with each device_event to complete. |
| | End of table |

4.10.1.8 sub_group class

The sub_group class encapsulates all functionality required to represent a particular sub-group within a parallel execution. It is not user-constructible.

The SYCL sub_group class provides the common by-value semantics (see Section 4.5.4).

A synopsis of the SYCL sub_group class is provided below. The member functions of the SYCL sub_group class are listed in Table 4.86. The additional common special member functions and common member functions are listed in 4.5.4 in Tables 4.3 and 4.4 respectively.

```
1 namespace sycl {
    class sub_group {
 2
 3 public:
 4
      using id_type = id<1>;
 5
 6
      using range_type = range<1>;
      using linear_id_type = uint32_t;
 7
 8
      static constexpr int dimensions = 1;
 9
      static constexpr memory_scope fence_scope = memory_scope::sub_group;
10
       /* -- common interface members -- */
11
12
13
      id<1> get_group_id() const;
14
15
      id<1> get_local_id() const;
16
17
      range<1> get_local_range() const;
18
19
      range<1> get_group_range() const;
20
21
      range<1> get_max_local_range() const;
22
23
      uint32_t get_group_linear_id() const;
24
25
      uint32_t get_local_linear_id() const;
26
```

```
27 uint32_t get_group_linear_range() const;
28
29 uint32_t get_local_linear_range() const;
30
31 bool leader() const;
32
33 };
34 } // sycl
```

| Member function | Description |
|--|--|
| <pre>id<1> get_group_id()const</pre> | Return an id representing the index of the |
| | sub-group within the work-group. |
| <pre>id<1> get_local_id()const</pre> | Return a SYCL id representing the calling |
| | work-item's position within the sub-group. |
| <pre>range<1> get_local_range()const</pre> | Return a SYCL range representing the size |
| | of the sub-group. This size may have been |
| | chosen by the programmer via an attribute, |
| | or chosen by the device compiler. |
| <pre>range<1> get_group_range()const</pre> | Return a range representing the number of |
| | sub-groups in the work-group. |
| <pre>range<1> get_max_local_range()const</pre> | Return a range representing the maximum |
| | number of work-items in any sub-group in |
| | the work-group. |
| <pre>uint32_t get_group_linear_id()const</pre> | Equivalent to get_group_id(). |
| <pre>uint32_t get_group_linear_range()const</pre> | Equivalent to get_group_range(). |
| <pre>uint32_t get_local_linear_id()const</pre> | Equivalent to get_local_id(). |
| <pre>uint32_t get_local_linear_range()const</pre> | Equivalent to get_local_range(). |
| <pre>bool leader()const</pre> | Return true for exactly one work-item in |
| | the sub-group, if the calling work-item is |
| | the leader of the sub-group, and false for all |
| | other work-items in the sub-group. |
| | The leader of the sub-group is determined |
| | during construction of the sub-group, and is |
| | invariant for the lifetime of the sub-group. |
| | The leader of the sub-group is guaranteed to |
| | be the work-item with a local id of 0. |
| | End of table |

4.10.2 Reduction variables

[Note for this provisional version: The reduction features described in this section support two alternative approaches for creating reducer objects and launching reduction kernels. Both alternatives are shown here to encourage feedback from implementers and developers, and it is expected that the final version of the SYCL specification will include only one approach. Please provide feedback on your preference or issues with either approach, by creating a new issue at https://github.com/KhronosGroup/SYCL-Docs/issues or extending an existing one. — end note]

All functionality related to reductions is captured by the reducer class, the reduction function, and the

parallel_reduce function.

The examples below demonstrate how to write a reduction kernel that performs two reductions simultaneously on the same input values, computing both the sum of all values in a buffer and the maximum value in the buffer.

In the first example, a reducer is created explicitly for each reduction variable and captured by the lambda passed to parallel_for.

```
1 buffer<int> valuesBuf { 1024 };
2 {
3
     // Initialize buffer on the host with 0, 1, 2, 3, ..., 1023
 4
     host_accessor a { valuesBuf };
5
     std::iota(a.begin(), a.end(), 0);
6 }
7
8
    // Buffers with just 1 element to get the reduction results
9
    int sumResult = 0;
10 buffer<int> sumBuf { &sumResult, 1 };
11 int maxResult = 0;
12 buffer<int> maxBuf { &maxResult, 1 };
13
14 myQueue.submit([&](handler& cgh) {
15
16
     // Input values to reductions are standard accessors
17
      auto inputValues = valuesBuf.get_access<access::mode::read>(cgh);
18
19
     // Create a reducer explicitly for each variable with reduction semantics
20
      auto sum = reducer(sumBuf.get_access<access::mode::read_write>(cgh), plus<>());
21
      auto max = reducer(maxBuf.get_access<access::mode::read_write>(cgh), maximum<>());
22
23
      // parallel_for operates on two reducers captured directly by the lambda
24
      cgh.parallel_for(range<1>{1024},
25
        [=](id<1> idx) {
26
          // plus<>() corresponds to += operator, so sum can be updated via += or combine()
27
          sum += inputValues[idx];
28
29
          // maximum<>() has no shorthand operator, so max can only be updated via combine()
30
          max.combine(inputValues[idx]);
31
     });
32 });
33
34
   // sumBuf and maxBuf contain the reduction results once the kernel completes
35
   assert(maxBuf.get_host_access()[0] == 1023 && sumBuf.get_host_access()[0] == 523776);
```

In the second example, one reducer is created explicitly and captured by the kernel lambda, and the other is created using the reduction function and passed to parallel_reduce as an argument.

```
1 buffer<int> valuesBuf { 1024 };
2 {
3     // Initialize buffer on the host with 0, 1, 2, 3, ..., 1023
4     host_accessor a { valuesBuf };
5     std::iota(a.begin(), a.end(), 0);
6 }
7
```

```
8 // Buffers with just 1 element to get the reduction results
 9 int sumResult = 0:
10 buffer<int> sumBuf { &sumResult, 1 };
11 int maxResult = 0;
12 buffer<int> maxBuf { &maxResult, 1 };
13
14
    myQueue.submit([&](handler& cgh) {
15
      // Input values to reductions are standard accessors
16
      auto inputValues = valuesBuf.get_access<access::mode::read>(cgh);
17
18
19
      // Create a reducer explicitly for a variable with reduction semantics
20
      auto sum = reducer(sumBuf.get_access<access::mode::read_write>(cgh), plus<>());
21
22
      // Create a temporary object describing a variable with reduction semantics
23
      auto maxReduction = reduction(maxBuf.get_access<access::mode::read_write>(cgh), maximum<>());
24
25
      // parallel_reduce operates on two reducers:
26
      // - sum is captured directly by the lambda
27
      // - max is created from maxReduction and passed to the lambda call operator
28
      cgh.parallel_reduce(range<1>{1024},
29
        maxReduction,
        [=](id<1> idx, auto& max) {
30
31
          // plus<>() corresponds to += operator, so sum can be updated via += or combine()
32
          sum += inputValues[idx];
33
34
          // maximum<>() has no shorthand operator, so max can only be updated via combine()
35
          max.combine(inputValues[idx]);
36
      });
37 });
38
39 // sumBuf and maxBuf contain the reduction results once the kernel completes
40
    assert(maxBuf.get_host_access()[0] == 1023 && sumBuf.get_host_access()[0] == 523776);
```

Reductions are supported for all trivially copyable types. If the reduction operator is non-associative or noncommutative, the behavior of a reduction may be non-deterministic. If multiple reductions reference the same reduction variable, or a reduction variable is accessed directly during the lifetime of a reduction (e.g. via an accessor or USM pointer), the behavior is undefined.

For user-defined reduction operators, an implementation should issue a compile-time warning if the functor does not contain a static constexpr member called identity_value, an identity is not specified in the call to reduction, and this is known to negatively impact performance (e.g. as a result of the implementation choosing a different reduction algorithm). For standard binary operations (e.g. plus) on arithmetic types, the implementation must determine the correct identity automatically in order to avoid performance penalties.

A reduction operation associated with a multi-dimensional accessor or a span represents an array reduction. An array reduction of size N is functionally equivalent to specifying N independent scalar reductions. The combination operations performed by an array reduction are limited to the accessible region of a buffer described by an accessor or the extent of a USM allocation described by a span, and access to elements outside of these regions results in undefined behavior.

4.10.2.1 reduction interface

The reduction interface is used to attach reduction semantics to a variable, by specifying: the reduction variable, the reduction operator and an optional identity value associated with the operator. The overloads of the interface are described in Table 4.87. The return value of the reduction interface is an implementation-defined object of unspecified type, which is interpreted by parallel_reduce to construct an appropriate reducer type as detailed in Section 4.10.2.2.

```
1
   template <typename AccessorT, typename BinaryOperation>
    __unspecified__ reduction(AccessorT vars, BinaryOperation combiner);
2
3
4
    template <typename AccessorT, typename BinaryOperation>
5
    __unspecified__ reduction(AccessorT vars, const T& identity, BinaryOperation combiner);
6
7
    template <typename T, typename BinaryOperation>
8
    __unspecified__ reduction(T* var, BinaryOperation combiner);
9
   template <typename T, typename BinaryOperation>
10
    __unspecified__ reduction(T* var, T& identity, BinaryOperation combiner);
11
12
13
   template <typename T, typename Extent, typename BinaryOperation>
14
    __unspecified__ reduction(span<T, Extent> vars, BinaryOperation combiner);
15
16 template <typename T, typename Extent, typename BinaryOperation>
    __unspecified__ reduction(span<T, Extent> vars, const T& identity, BinaryOperation combiner);
17
```

| Function | Description |
|---|--|
| Function | Description |
| <pre>reduction<accessort, binaryoperation="">(AccessorT vars</accessort,></pre> | Construct an unspecified object represent- |
| , BinaryOperation combiner) | ing a reduction of the variable(s) described |
| | by vars using the combination operation |
| | specified by combiner. If the access mode |
| | of vars is <pre>access::mode::read_write</pre> |
| | then the reduction operation includes the |
| | original value(s) of the variable(s) described |
| | by vars. vars must not be a placeholder |
| | accessor. |
| | |
| | Available only when: (accessMode |
| | == access::mode::read write |
| | <pre> accessMode == access::mode::</pre> |
| | discard write)&& accessTarget == |
| | access thereat is a lobal buffer |
| | accesstargetgiobal_builer |
| | Continued on next page |

Table 4.87: Overloads of the reduction interface.

| Function | Description |
|---|--|
| <pre>reduction<accessort, binaryoperation="">(AccessorT vars</accessort,></pre> | Construct an unspecified object represent- |
| <pre>, AccessorT::value_type identity, BinaryOperation</pre> | ing a reduction of the variable(s) described |
| combiner) | by vars using the combination operation |
| | specified by combiner. The value of |
| | identity may be used by the implemen- |
| | tation to initialize temporary accumulation |
| | variables; using an identity value that is |
| | not the identity value of the combination |
| | operation specified by combiner results in |
| | undefined benavior. If the access mode |
| | of vars is access::mode::read_write |
| | original value(s) of the variable(s) described |
| | by yors wars must not be a placeholder |
| | accessor |
| | |
| | Available only when: (accessMode |
| | == access::mode::read write |
| | <pre> accessMode == access::mode::</pre> |
| | discard_write)&& accessTarget == |
| | <pre>access::target::global_buffer</pre> |
| <pre>reduction<t, binaryoperation="">(T* var,</t,></pre> | Construct an unspecified object represent- |
| BinaryOperation combiner) | ing a reduction of the variable described by |
| | var using the combination operation speci- |
| | fied by combiner. The reduction operation |
| | includes the original value of the variable |
| | described by var. |
| <pre>reduction<t, binaryoperation="">(T* var, T identity,</t,></pre> | Construct an unspecified object represent- |
| BinaryOperation combiner) | ing a reduction of the variable described by |
| | var using the combination operation spec- |
| | ified by combiner. The value of identity |
| | may be used by the implementation to |
| | using an identity value that is not the iden |
| | tity value of the combination operation spec |
| | ified by combiner results in undefined be- |
| | havior. The reduction operation includes the |
| | original value of the variable described by |
| | var. |
| <pre>reduction<t, binaryoperation="">(span<t, extent=""> vars,</t,></t,></pre> | Construct an unspecified object represent- |
| BinaryOperation combiner) | ing a reduction of the variable(s) described |
| | by vars using the combination operation |
| | specified by combiner. The reduction op- |
| | eration includes the original value(s) of the |
| | variable described by vars. |
| | Continued on next page |

Table 4.87: Overloads of the **reduction** interface.

| Function | Description |
|--|---|
| <pre>reduction<t, binaryoperation="">(span<t, extent=""> vars,</t,></t,></pre> | Construct an unspecified object repre- |
| T identity, BinaryOperation combiner) | senting a reduction of the variable(s) de- scribed by vars using the combination op- eration specified by combiner. The value of identity may be used by the implemen- tation to initialize temporary accumulation variables; using an identity value that is not the identity value of the combination op- eration specified by combiner results in un- |
| | defined behavior. The reduction operation includes the original value(s) of the vari- |
| | able(s) described by vars. |
| | End of table |

Table 4.87: Overloads of the reduction interface.

4.10.2.2 reducer class

The **reducer** class defines the interface between a work-item and a reduction variable during the execution of a SYCL kernel, restricting access to the underlying reduction variable. The intermediate values of a reduction variable cannot be inspected during kernel execution, and the variable cannot be updated using anything other than the reduction's specified combination operation. The combination order of different reducers is unspecified, as are when and how the value of each reducer is combined with the original reduction variable.

A reducer can be constructed explicitly by a user and bound to a command group handler, or constructed by an implementation of parallel_reduce given the return value of a call to the reduction function. To enable compiletime specialization of reduction algorithms, the implementation of the reducer class is unspecified, except for the functions and operators defined in Tables 4.88, 4.89 and 4.90. It is recommended that developers use auto in place of specifying the template arguments of a reducer directly.

An implementation must guarantee that it is safe for each concurrently executing work-item in a kernel to call the combine function of a **reducer** in parallel. An implementation is free to re-use reducer variables (e.g. across work-groups scheduled to the same compute unit) if it can guarantee that it is safe to do so.

The member functions of the **reducer** class are listed in Table 4.89. Additional shorthand operators may be made available for certain combinations of reduction variable type and combination operation, as described in Table 4.90.

```
1 // Exposition only
2
    template <typename T, typename BinaryOperation, int Dimensions>
3
   class reducer {
4
5
      template <access_mode Mode>
6
      reducer(accessor<T, Dimensions, Mode> vars, BinaryOperation combiner);
 7
8
      template <access_mode Mode>
9
      reducer(accessor<T, Dimensions, Mode> vars, const T& identity, BinaryOperation combiner);
10
11
      reducer(T* var, BinaryOperation combiner, handler& cgh);
12
13
      reducer(T* var, const T& identity, BinaryOperation combiner, handler& cgh);
```

```
14
15
      template <typename Extent>
16
      reducer(span<T, Extent> vars, BinaryOperation combiner, handler& cgh);
17
18
      template <typename Extent>
19
      reducer(span<T, Extent> vars, const T& identity, BinaryOperation combiner, handler& cgh);
20
21
      reducer(const reducer<T.BinaryOperation.Dimensions>&) = delete;
22
      reducer<T,BinaryOperation,Dimensions>& operator(const reducer<T,BinaryOperation,Dimensions>&) =
          delete;
23
24
      /* Only available if Dimensions == 0 */
25
     void combine(const T& partial);
26
27
      /* Only available if Dimensions > 1 */
28
      __unspecified__ &operator[](size_t index) const;
29
30
      /* Only available if identity value is known */
31
     T identity() const;
32
33 };
34
35 template <typename T>
36 void operator+=(reducer<T,plus<T>,0>&, const T&);
37
38 template <typename T>
39
   void operator*=(reducer<T,multiplies<T>,0>&, const T&);
40
41 /* Only available for integral types */
42 template <typename T>
43 void operator&=(reducer<T,bit_and<T>,0>&, const T&);
44
45 /* Only available for integral types */
46 template <typename T>
47 void operator |= (reducer<T, bit_or<T>, 0>&, const T&);
48
49~ /* Only available for integral types */
50 template <typename T>
51 void operator^=(reducer<T,bit_xor<T>,0>&, const T&);
52
53 /* Only available for integral types */
54 template <typename T>
```

55 void operator++(reducer<T,plus<T>,0>&);

| Constructor | Description |
|--|---|
| <pre>reducer<access_mode mode="">(accessor<t, dimensions,<br="">Mode> vars, BinaryOperation combiner)</t,></access_mode></pre> | Construct a reducer representing a reduc- tion of the variable(s) described by vars using the combination operation specified by combiner. If the access mode of vars is access::mode::read_write then the reduc- tion operation includes the original value(s) of the variable(s) described by vars. vars must not be a placeholder accessor. |
| | <pre>Available only when: (accessMode == access::mode::read_write accessMode == access::mode:: discard_write)&& accessTarget == access::target::global_buffer</pre> |
| <pre>reducer<access_mode mode="">(accessor<t, dimensions,<br="">Mode> vars, const T& identity, BinaryOperation combiner)</t,></access_mode></pre> | Construct a reducer representing a reduc- tion of the variable(s) described by vars using the combination operation specified by combiner. The value of identity may be used by the implementation to initialize temporary accumulation variables; using an identity value that is not the identity value of the combination operation specified by combiner results in undefined behavior. If the access mode of vars is access::mode ::read_write then the reduction operation includes the original value(s) of the vari- able(s) described by vars. vars must not be a placeholder accessor. Available only when: (accessMode == access::mode::read_write accessMode == access::mode:: |
| <pre>reducer(T* var, BinaryOperation combiner, handler& cgh)</pre> | <pre>discard_write)&& accessTarget == access::target::global_buffer Construct a reducer representing a reduc- tion of the variable described by var us-</pre> |
| | ing the combination operation specified by combiner. The reduction operation includes the original value of the variable described by var. |
| | Continued on next page |

 Table 4.88: Constructors of the reducer class.

| Constructor | Description |
|--|--|
| <pre>reducer(T* var, const T& identity, BinaryOperation</pre> | Construct a reducer representing a reduc- |
| combiner, handler& cgh) | tion of the variable described by var us- |
| | ing the combination operation specified by |
| | combiner. The value of identity may |
| | be used by the implementation to initial- |
| | ize temporary accumulation variables; using |
| | an identity value that is not the identity |
| | value of the combination operation specified |
| | by combiner results in undefined behavior. |
| | The reduction operation includes the origi- |
| | nal value of the variable described by var. |
| <pre>reducer<extent>(span<t, extent=""> vars,</t,></extent></pre> | Construct a reducer representing a reduc- |
| BinaryOperation combine, handler& cgh) | tion of the variable(s) described by vars |
| | using the combination operation specified |
| | by combiner. The reduction operation in- |
| | cludes the original value(s) of the variable |
| | described by vars. |
| <pre>reducer<extent>(span<t, extent=""> vars, const T&</t,></extent></pre> | Construct a reducer representing a reduc- |
| identity, BinaryOperation combiner, handler& cgh) | tion of the variable(s) described by vars |
| | using the combination operation specified |
| | by combiner. The value of identity may |
| | be used by the implementation to initial- |
| | ize temporary accumulation variables; using |
| | an identity value that is not the identity |
| | value of the combination operation specified |
| | by combiner results in undefined behavior. |
| | The reduction operation includes the origi- |
| | nal value(s) of the variable(s) described by |
| | vars. |
| | End of table |

Table 4.88: Constructors of the reducer class.

| Member function | Description |
|--|--|
| <pre>void combine(const T& partial)const</pre> | Combine the value of partial with the reduction variable associated with this reducer. |
| | Continued on next page |

Table 4.89: Member functions of the reducer class.

| Member function | Description |
|---|--|
| <pre>unspecified &operator[](size_t index)const</pre> | Available only when: Dimensions > 1. Re- |
| | turns an instance of an undefined intermedi- |
| | ate type representing a reducer of the same |
| | type as this reducer, with the dimension- |
| | ality Dimensions-1 and containing an im- |
| | plicit SYCL id with index Dimensions set to |
| | index. The intermediate type returned must |
| | provide all member functions and operators |
| | defined by the reducer class that are appro- |
| | priate for the type it represents (including |
| | this subscript operator). |
| T identity()const | Return the identity value of the combina- |
| | tion operation associated with this reducer. |
| | Only available if the identity value is known |
| | to the implementation, or was specified ex- |
| | plicitly in the call to reduction that returned |
| | this reducer. |
| | End of table |

Table 4.89: Member functions of the reducer class.

| Operator | Description |
|---|---|
| template <typename t=""></typename> | Equivalent to calling accum.combine(|
| <pre>void operator+=(reducer<t,plus<t>,0>& accum,</t,plus<t></pre> | partial). |
| const T& partial) | |
| template <typename t=""></typename> | Equivalent to calling accum.combine(|
| <pre>void operator*=(reducer<t,multiplies<t>,0>&</t,multiplies<t></pre> | partial). |
| accum, const T& partial) | |
| template <typename t=""></typename> | Equivalent to calling accum.combine(|
| <pre>void operator =(reducer<t,bit_or<t>,0>& accum,</t,bit_or<t></pre> | partial). Only available for integral |
| <pre>const T& partial)</pre> | types. |
| template <typename t=""></typename> | Equivalent to calling accum.combine(|
| <pre>void operator&=(reducer<t,bit_and<t>,0>& accum,</t,bit_and<t></pre> | partial). Only available for integral |
| <pre>const T& partial)</pre> | types. |
| template <typename t=""></typename> | Equivalent to calling accum.combine(|
| <pre>void operator^=(reducer<t,bit_xor<t>,0>& accum,</t,bit_xor<t></pre> | partial). Only available for integral |
| const T& partial) | types. |
| template <typename t=""></typename> | Equivalent to calling accum.combine(1). |
| <pre>void operator++(reducer<t,plus<t>,0>& accum)</t,plus<t></pre> | Only available for integral types. |
| | End of table |

Table 4.90: Operators of the reducer class.

4.10.3 Command group scope

A command group scope, as defined in Section 3.6.1, may execute a single command such as invoking a kernel, copying memory, or executing a host task. It is legal for a command group scope to statically contain more than one call to a command function, but any single execution of the command group function object may execute

no more than one command. The statements that call commands together with the statements that define the requirements for a kernel form the command group function object. The command group function object takes as a parameter an instance of the command group handler class which encapsulates all the member functions executed in the command group scope. The member functions and objects defined in this scope will define the requirements for the kernel execution or explicit memory operation, and will be used by the SYCL runtime to evaluate if the operation is ready for execution. Host code within a command group function object (typically setting up requirements) is executed once, before the command group submit call returns. This abstraction of the kernel execution unifies the data with its processing, and consequently allows more abstraction and flexibility in the parallel programming models that can be implemented on top of SYCL.

The command group function object and the handler class serve as an interface for the encapsulation of command group scope. A SYCL kernel function is defined as a function object. All the device data accesses are defined inside this group and any transfers are managed by the SYCL runtime. The rules for the data transfers regarding device and host data accesses are better described in the data management section (4.7), where buffers (4.7.2) and accessor (4.7.6) classes are described. The overall memory model of the SYCL application is described in Section 3.7.1.

It is possible to obtain events for the start of the command group function object, the kernel starting, and the command group completing. These events are most useful for profiling, because safe synchronization in SYCL requires synchronization on buffer availability, not on kernel completion. This is because the memory that data is stored in upon kernel completion is not rigidly specified. The events are provided at the submission of the command group function object to the queue to be executed on.

It is possible for a command group function object to fail to enqueue to a queue, or for it to fail to execute correctly. A user can therefore supply a secondary queue when submitting a command group to the primary queue. If the SYCL runtime fails to enqueue or execute a command group on a primary queue, it can attempt to run the command group on the secondary queue. The circumstances in which it is, or is not, possible for a SYCL runtime to fall-back from primary to secondary queue are unspecified in the specification. Even if a command group is run on the secondary queue, the requirement that host code within the command group is executed exactly once remains, regardless of whether the fallback queue is used for execution.

The command group handler class provides the interface for all of the member functions that are able to be executed inside the command group scope, and it is also provided as a scoped object to all of the data access requests. The command group handler class provides the interface in which every command in the command group scope will be submitted to a queue.

4.10.4 Command group handler class

A command group handler object can only be constructed by the SYCL runtime. All of the accessors defined in command group scope take as a parameter an instance of the command group handler, and all the kernel invocation functions are member functions of this class.

The constructors of the SYCL handler class are described in Table 4.91.

It is disallowed for an instance of the SYCL handler class to be moved or copied.

```
1 namespace sycl {
2
3 class handler {
4 private:
5
6 // implementation defined constructor
```

```
7
     handler(___unspecified___);
8
9
     public:
10
11
      template <typename dataT, int dimensions, access::mode accessMode,
12
        access::target accessTarget, access::placeholder isPlaceholder>
13
      void require(accessor<dataT, dimensions, accessMode, accessTarget,</pre>
14
                   isPlaceholder> acc);
15
16
      //---- Backend interoperability interface
17
      11
18
      template <typename T>
19
      void set_arg(int argIndex, T && arg);
20
21
      template <typename... Ts>
22
      void set_args(Ts &&... args);
23
     //----- Kernel dispatch API
24
25
     11
26
     // Note: In all kernel dispatch functions, the template parameter
27
     // "typename kernelName" is optional.
28
      11
29
      template <typename KernelName, typename KernelType>
30
      void single_task(const KernelType &kernelFunc);
31
32
      template <typename KernelName, typename KernelType, int dimensions>
33
      void parallel_for(range<dimensions> numWorkItems, const KernelType &kernelFunc);
34
35
      template <typename KernelName, typename KernelType, int dimensions>
36
      void parallel_for(range<dimensions> numWorkItems,
                        id<dimensions> workItemOffset, const KernelType &kernelFunc);
37
38
39
      template <typename KernelName, typename KernelType, int dimensions>
40
      void parallel_for(nd_range<dimensions> executionRange, const KernelType &kernelFunc);
41
42
      template <typename KernelName, typename WorkgroupFunctionType, int dimensions>
43
      void parallel_for_work_group(range<dimensions> numWorkGroups,
44
                                   const WorkgroupFunctionType &kernelFunc);
45
46
      template <typename KernelName, typename WorkgroupFunctionType, int dimensions>
47
      void parallel_for_work_group(range<dimensions> numWorkGroups,
48
                                   range<dimensions> workGroupSize,
49
                                   const WorkgroupFunctionType &kernelFunc);
50
51
      void single_task(kernel syclKernel);
52
53
      template <int dimensions>
54
      void parallel_for(range<dimensions> numWorkItems, kernel syclKernel);
55
56
      template <int dimensions>
57
      void parallel_for(range<dimensions> numWorkItems,
58
                        id<dimensions> workItemOffset, kernel syclKernel);
59
60
      template <int dimensions>
61
      void parallel_for(nd_range<dimensions> ndRange, kernel syclKernel);
```

| 62 | |
|------------|--|
| 63 | // Explicit memory operation APIs |
| 64 | |
| 65 | <pre>template <typename access::="" access::mode="" access::target="" dim_src,="" int="" isplaceholder,<="" mode_src,="" placeholder="" pre="" t_src,="" tgt_src,=""></typename></pre> |
| 66 | typename T_dest> |
| 67 68 | <pre>void copy(accessor<t_src, dim_src,="" isplaceholder="" mode_src,="" tgt_src,=""> src,</t_src,></pre> |
| 69 | Star.Sharea_ptr (1_acst) acst); |
| 70 | template <typename src.<="" t="" td=""></typename> |
| 71 | <pre>typename T_dest, int dim_dest, access::mode mode_dest, access::target tgt_dest, access ::placeholder isPlaceholder></pre> |
| 72 | <pre>void copy(std::shared_ptr<t_src> src,</t_src></pre> |
| 73 | <pre>accessor<t_dest, dim_dest,="" isplaceholder="" mode_dest,="" tgt_dest,=""> dest);</t_dest,></pre> |
| 74 | |
| 75 76 | <pre>template <typename access::="" access::mode="" access::target="" dim_src,="" int="" isplaceholder,="" mode_src,="" placeholder="" t_dect="" t_src,="" tgt_src,="" tupename=""></typename></pre> |
| 70 | upid conv(accessor T are dim are mode are tat are isPlaceholder) are |
| 78 | T doct *doct). |
| 79 | I_dest dest), |
| 80 | template <typename src.<="" t="" td=""></typename> |
| 81 | typename T_dest, int dim_dest, access::mode mode_dest, access::target tgt_dest, access |
| | ::placeholder isPlaceholder> |
| 82 | <pre>void copy(const T_src *src,</pre> |
| 83 | <pre>accessor<t_dest, dim_dest,="" isplaceholder="" mode_dest,="" tgt_dest,=""> dest);</t_dest,></pre> |
| 84 | |
| 85 | <pre>template <typename access::="" access::mode="" access::target="" dim_src,="" int="" isplaceholder_src,<="" mode_src,="" placeholder="" pre="" t_src,="" tgt_src,=""></typename></pre> |
| 86 | <pre>typename T_dest, int dim_dest, access::mode mode_dest, access::target tgt_dest, access ::placeholder isPlaceholder_dest></pre> |
| 87 | <pre>void copy(accessor<t_src, dim_src,="" isplaceholder_src="" mode_src,="" tgt_src,=""> src,</t_src,></pre> |
| 88 89 | <pre>accessor<t_dest, dim_dest,="" isplaceholder_dest="" mode_dest,="" tgt_dest,=""> dest);</t_dest,></pre> |
| 90 | <pre>template <typename access::mode="" access::placeholder="" access::target="" dim,="" int="" isplaceholder="" mode,="" t,="" tgt,=""></typename></pre> |
| 91 92 | <pre>void update_host(accessor<t, dim,="" isplaceholder="" mode,="" tgt,=""> acc);</t,></pre> |
| 93 | <pre>template <typename access::mode="" access::placeholder="" access::target="" dim,="" int="" isplaceholder="" mode,="" t,="" tgt,=""></typename></pre> |
| 94 05 | <pre>void fill(accessor<t, dim,="" isplaceholder="" mode,="" tgt,=""> dest, const T& src);</t,></pre> |
| 96 97 | <pre>void use_module(const module_state::executable> &execModule);</pre> |
| 98 | template <typename t=""></typename> |
| 99 | <pre>void use_module(const module<module_state::executable> &execModule</module_state::executable></pre> |
| 100 101 | T deviceImageSelector); |
| 102 | <pre>template<auto& s=""></auto&></pre> |
| 103 | <pre>bool has_specialization_constant() const noexcept;</pre> |
| 104 | |
| 105 | <pre>template<auto& s=""></auto&></pre> |
| 106 | <pre>typename std::remove_reference_t<decltype(s)>::type get_specialization_constant();</decltype(s)></pre> |
| 107 | |
| 108 | }; |
| | |

109 } // namespace sycl

| Constructor | Description |
|---------------------------------|---|
| <pre>handler(unspecified)</pre> | Unspecified implementation defined con- |
| | structor. |
| | End of table |

Table 4.91: Constructors of the handler class.

4.10.5 Class kernel_handler

Functionality and queries that are unique to the invocation of a SYCL kernel function are made available at the kernel scope via a kernel handler. A kernel handler is associated with the SYCL kernel function that is being invoked and the module and device image used by the SYCL runtime. A kernel handler is represented by the class kernel_handler.

The kernel handler is optional, and is only used if the SYCL kernel function has a kernel_handler as an additional parameter, in which case the SYCL runtime will construct an instance of kernel_handler and pass it to the SYCL kernel function as an argument. A kernel_handler is not user constructible and can only be constructed by the SYCL runtime.

1 class kernel_handler;

4.10.5.1 Constructors

- 1 kernel_handler(__unspecified__); // (1)
 - 1. Effects: Unspecified private constructor for the SYCL runtime to use to construct a kernel handler.

4.10.5.2 Member functions

```
1 template<auto& S>
```

- 2 bool has_specialization_constant() const noexcept; // (1)
 - 1. *Returns:* true if any of the SYCL kernel functions represented by the associated module contains the specialization constant represented by the specialization_id at the address S, otherwise returns false.

```
1 template<auto& S>
```

- 2 typename std::remove_reference_t<decltype(S)>::type get_specialization_constant(); // (1)
 - 1. *Returns:* The value of the specialization constant associated with the specialization_id at the address S, from the associated module, if the specialization constant has been set, otherwise returns the default value.

Effects: If the associated module is associated with a host context or this->has_specialization_constant <S>() evaluates to false this member function is undefined.

4.10.6 SYCL functions for adding requirements

When an accessor is created from a command group handler, a **requirement** is implicitly added to the command group for the accessor's data. However, this does not happen when creating a *placeholder* accessor. In order to create a **requirement** for a *placeholder* accessor, code must call the handler::require() member function.

| Member function | Description |
|---|--|
| <pre>template <typename datat,="" dimensions,<="" int="" pre=""></typename></pre> | Requires access to the memory object asso- |
| <pre>access::mode accessMode, access::target accessTarget</pre> | ciated with the accessor. |
| <pre>, access::placeholder isPlaceholder></pre> | The command group now has a require- |
| <pre>void require(accessor<datat, dimensions,<="" pre=""></datat,></pre> | ment to gain access to the given memory |
| <pre>accessMode, accessTarget, isPlaceholder></pre> | object before executing the kernel. If the ac- |
| acc) | cessor has already been registered with the |
| | command group, calling this function has no |
| | effect. |
| | Throws exception with the errc:: |
| | accessor_error error code if (acc.empty |
| | ()== true). |
| | End of table |

Table 4.92: Member functions of the handler class.

4.10.7 SYCL functions for invoking kernels

Kernels can be invoked as *single tasks*, basic *data-parallel* kernels, nd-range in work-groups, or *hierarchical parallelism*.

Each function takes an optional kernel name template parameter. The user may optionally provide a kernel name, otherwise an implementation defined name will be generated for the kernel.

All the functions for invoking kernels are member functions of the command group handler class 4.10.4, which is used to encapsulate all the member functions provided in a command group scope. Table 4.93 lists all the members of the handler class related to the kernel invocation.

| Member function | Description |
|--|--|
| template <typename t=""></typename> | Set a kernel argument for a kernel through |
| <pre>void set_arg(int argIndex, T &&arg)</pre> | the SYCL backend interoperability inter- |
| | face. The index value specifies which pa- |
| | rameter of the low-level kernel is being set |
| | and arg specifies the kernel argument. |
| | Index 0 is the first parameter. |
| | The argument can be either a SYCL acces- |
| | sor, a SYCL sampler or a trivially copyable |
| | C++ type. |
| | Note it is invalid to set arguments to SYCL |
| | kernel function objects. |
| | Continued on next page |

Table 4.93: Member functions of the handler class.

| Member function | Description |
|---|---|
| <pre>template <typename ts=""></typename></pre> | Set all the given kernel args arguments for |
| <pre>void set_args(Ts && args)</pre> | an kernel through the SYCL backend inter- |
| | operability interface, as if set_arg() was |
| | used with each of them in the same order |
| | and increasing index always starting at 0. |
| <pre>template <typename kernelname,="" kerneltype="" typename=""></typename></pre> | Defines and invokes a SYCL kernel func- |
| <pre>void single_task(const KernelType &kernelFunc)</pre> | tion as a lambda function or a named |
| | function object type. Specification of |
| | a kernel name (typename KernelName), as |
| | described in Section 4.10.7, is optional. |
| | take a harmal handlan in which acce the |
| | take a kernel_nandler in which case the |
| | sice functime will construct an instance of |
| | Defines and involves a SVCL formal function |
| int dimensions | as a lambda function or a named function |
| woid parallel for | object type for the specified range and given |
| range <dimensions> numWorkItems</dimensions> | an item or integral type (e g int size t) if |
| const KernelType &kernelFunc) | range is 1-dimensional, for indexing in the |
| | indexing space defined by range. Generic |
| | kernel functions are permitted, in that case |
| | the argument type is an item. Specification |
| | of a kernel name (typename KernelName), |
| | as described in Section 4.10.7, is optional. |
| | The callable KernelType can optionally take |
| | a kernel_handler as it's last parameter, in |
| | which case the SYCL runtime will construct |
| | an instance of kernel_handler and pass it to |
| | KernelType. |
| <pre>template <typename kernelname,="" kerneltype,<="" pre="" typename=""></typename></pre> | Defines and invokes a SYCL kernel func- |
| int dimensions> | tion as a lambda function or a named func- |
| void parallel_for(| tion object type, for the specified range and |
| range <dimensions> numWorkItems,</dimensions> | offset and given an item or integral type |
| 1d <dimensions> workitemoffset,</dimensions> | (e.g int, size_t), il range is 1-dimensional, |
| const kernellype &kernelfunc) | fined by range Generic kernel functions |
| | are permitted in that case the argument |
| | type is an item Specification of a ker- |
| | nel name (typename KernelName), as de- |
| | scribed in Section 4.10.7. is optional. The |
| | callable KernelType can optionally take a |
| | kernel_handler as it's last parameter, in |
| | which case the SYCL runtime will construct |
| | an instance of kernel_handler and pass it to |
| | KernelType. |
| | Continued on next page |

| Member function | Description |
|---|--|
| <pre>template <typename kernelname,="" kerneltype,<="" pre="" typename=""></typename></pre> | Defines and invokes a SYCL kernel func- |
| int dimensions> | tion as a lambda function or a named |
| <pre>void parallel_for(</pre> | function object type, for the specified nd- |
| <pre>nd_range<dimensions> executionRange,</dimensions></pre> | range and given an nd-item for indexing |
| <pre>const KernelType &kernelFunc)</pre> | in the indexing space defined by the nd- |
| | range. Generic kernel functions are per- |
| | mitted, in that case the argument type |
| | is an nd-item. Specification of a ker- |
| | nel name (typename KernelName), as de- |
| | scribed in Section 4.10.7, is optional. The |
| | callable KernelType can optionally take a |
| | kernel_handler as it's last parameter, in |
| | which case the SYCL runtime will construct |
| | an instance of kernel_handler and pass it to |
| | Kernellype. |
| template <typename kernelname,="" td="" typename<=""><td>Defines and invokes a hierarchical kernel as</td></typename> | Defines and invokes a hierarchical kernel as |
| workgroupFunctionType, Int dimensions> | work group to lough. Generic kernel func |
| void parallel_toi_work_group(| tions are permitted in that case the argument |
| const WorkgroupFunctionType &kernelFunc) | type is a group. May contain multiple calls |
| const workgrouprunctionrype akernerruncy | to parallel for work item() member |
| | functions representing the execution on each |
| | work-item. Launches num work groups |
| | work-groups of runtime-defined size. De- |
| | scribed in detail in 4.10.7. The callable |
| | WorkgroupFunctionType can optionally take |
| | a kernel_handler as it's last parameter, in |
| | which case the SYCL runtime will construct |
| | an instance of kernel_handler and pass it to |
| | WorkgroupFunctionType. |
| template <typename kernelname,="" td="" typename<=""><td>Defines and invokes a hierarchical kernel as</td></typename> | Defines and invokes a hierarchical kernel as |
| WorkgroupFunctionType, int dimensions> | a lambda function encoding the body of each |
| <pre>void parallel_for_work_group(</pre> | work-group to launch. Generic kernel func- |
| <pre>range<dimensions> numWorkGroups,</dimensions></pre> | tions are permitted, in that case the argument |
| <pre>range<dimensions> workGroupSize,</dimensions></pre> | type is a group. May contain multiple calls |
| <pre>const WorkgroupFunctionType &kernelFunc)</pre> | to parallel_for_work_item member func- |
| | tions representing the execution on each |
| | work-item. Launches num_work_groups |
| | items and Described in detail in 4.10.7 |
| | The colleble Workgroup Function Turne |
| | can optionally take a kornal handlor |
| | as it's last parameter in which case |
| | the SYCL runtime will construct an in- |
| | stance of kernel handler and pass it to |
| | WorkgroupFunctionTvpe. |
| | Continued on next page |

| Member function | Description |
|--|--|
| <pre>template <typename kernelname,="" kerneltype,<="" pre="" typename=""></typename></pre> | Defines and invokes a SYCL kernel function |
| <pre>int dimensions, typename Reductions></pre> | as a lambda function or a named function |
| <pre>void parallel_reduce(</pre> | object type, for the specified range and given |
| <pre>range<dimensions> numWorkItems, Reductions</dimensions></pre> | an id or item for indexing in the indexing |
| reductions, const KernelType &kernelFunc) | space defined by range. Specification of |
| | a kernel name (typename KernelName), as |
| | described in Section 4.10.7, is optional. The |
| | callable KernelType can optionally take a |
| | kernel_handler as it's last parameter, in |
| | which case the SYCL runtime will construct |
| | an instance of kernel_handler and pass it |
| | to KernelType. |
| | |
| | The reductions parameter pack con- |
| | sists of 1 or more objects created by the |
| | reduction function. For each object in |
| | reductions, the kernel functor should take |
| | an additional parameter corresponding to |
| | that object's reducer type. |
| template <typename kernelname,="" kerneltype,<="" td="" typename=""><td>Defines and invokes a SYCL kernel function</td></typename> | Defines and invokes a SYCL kernel function |
| int dimensions, typename Reductions> | as a lambda function or a named function |
| vold parallel_reduce(| object type, for the specified range and |
| range <dimensions> numWorkItems,</dimensions> | offset and given an id or item for indexing |
| 1d <dimensions> workItemOffset, Reductions</dimensions> | in the indexing space defined by range. |
| reductions, const KernelType &kernelFunc) | Specification of a kernel name (typename |
| | KernelName), as described in Section 4.10.7, |
| | is optional. The callable kernellype can |
| | optionally take a kernel_nandler as it's |
| | last parameter, in which case the SYCL |
| | fundine will construct an instance of |
| | kernel_nandler and pass it to kernellype. |
| | The reductions parameter pack con- |
| | sists of 1 or more objects created by the |
| | reduction function. For each object in |
| | reductions, the kernel functor should take |
| | an additional parameter corresponding to |
| | that object's reducer type. |
| | Continued on next page |

| Member function | Description |
|---|--|
| <pre>template <typename kernelname,="" kerneltype,<="" pre="" typename=""></typename></pre> | Defines and invokes a SYCL kernel function |
| <pre>int dimensions, typename Reductions></pre> | as a lambda function or a named function |
| <pre>void parallel_reduce(</pre> | object type, for the specified nd-range |
| <pre>nd_range<dimensions> executionRange, Reductions</dimensions></pre> | and given an nd-item for indexing in the |
| <pre> reductions, const KernelType &kernelFunc)</pre> | indexing space defined by the nd-range. Specification of a kernel name (typename KernelName), as described in Section 4.10.7, is optional. The callable KernelType can optionally take a kernel_handler as it's last parameter, in which case the SYCL runtime will construct an instance of kernel_handler and pass it to KernelType. as described in 4.10.7. |
| | The reductions parameter pack con- sists of 1 or more objects created by the reduction function. For each object in reductions, the kernel functor should take an additional parameter corresponding to that object's reducer type. |
| <pre>void single_task(kernel syclKernel)</pre> | Invokes a pre-compiled kernel which exe- cutes exactly once. |
| <pre>template <int dimensions=""> void parallel_for(</int></pre> | Invokes a pre-compiled kernel for the speci- |
| <pre>range<dimensions> numWorkItems,</dimensions></pre> | fied range and given an item or integral type |
| kernel syclKernel) | (e.g int, size_t), if range is 1-dimensional, for indexing in the indexing space defined by range. Generic kernel functions are per- mitted, in that case the argument type is an item. Described in detail in 4.10.7. |
| <pre>template <int dimensions=""> void parallel_for(</int></pre> | Invokes a pre-compiled kernel for the spec- |
| <pre>range<dimensions> numWorkItems,</dimensions></pre> | ified range and offset and given an item or |
| <pre>id<dimensions> workItemOffset, kernel syclKernel)</dimensions></pre> | integral type (e.g int, size_t), if range is 1-dimensional, for indexing in the indexing space defined by range. Generic kernel func- tions are permitted, in that case the argu- ment type is an item. Described in detail in 4.10.7. |
| <pre>template <int dimensions=""> void parallel_for(</int></pre> | Invokes a pre-compiled kernel for the spec- |
| <pre>nd_range<dimensions> ndRange,</dimensions></pre> | Iffied ndrange and given an nd_item for in- |
| kernel syclKernel) | dexing in the indexing space defined by the |
| | nd_range. Generic kernel functions are per- |
| | nd_item. Described in detail in 4.10.7. |
| | End of table |
4.10.7.1 single_task invoke

SYCL provides a simple interface to enqueue a kernel that will be sequentially executed on a device. Only one instance of the kernel will be executed. This interface is useful as a primitive for more complicated parallel algorithms, as it can easily create a chain of sequential tasks on a SYCL device with each of them managing its own data transfers.

This function can only be called inside a command group using the handler object created by the runtime. Any accessors that are used in a kernel should be defined inside the same command group.

Local accessors are disallowed for single task invocations.

For single tasks, the kernel member function takes no parameters, as there is no need for index space classes in a unary index space.

A kernel_handler can optionally be passed as a parameter to the SYCL kernel function that is invoked by single_task for the purpose explained in Section 4.10.5.

4.10.7.2 parallel_for invoke

The parallel_for member function of the SYCL handler class provides an interface to define and invoke a SYCL kernel function in a command group, to execute in parallel execution over a 3 dimensional index space. There are three overloads of the parallel_for member function which provide variations of this interface, each with a different level of complexity and providing a different set of features.

For the simplest case, users need only provide the global range (the total number of work-items in the index space) via a SYCL range parameter, and the SYCL runtime will select a local range (the number of work-items in each work-group). The local range chosen by the SYCL runtime is entirely implementation defined. In this case the function object that represents the SYCL kernel function must take one of: 1) a single SYCL item parameter, 2) single generic parameter (template parameter or auto), 3) any other type implicitly converted from SYCL item, representing the currently executing work-item within the range specified by the range parameter.

The execution of the kernel function is the same whether the parameter to the SYCL kernel function is a SYCL id or a SYCL item. What differs is the functionality that is available to the SYCL kernel function via the respective interfaces.

Below is an example of invoking a SYCL kernel function with parallel_for using a lambda function, and passing a SYCL id parameter. In this case only the global id is available. This variant of parallel_for is designed for

when it is not necessary to query the global range of the index space being executed across, or the local (workgroup) size chosen by the implementation.

```
1 myQueue.submit([&](handler & cgh) {
2 accessor acc { myBuffer, cgh, write_only };
3
4 cgh.parallel_for(range<1>(numWorkItems),
5 [=] (id<1> index) {
6 acc[index] = 42.0f;
7 });
8 });
```

Below is an example of invoking a SYCL kernel function with parallel_for using a lambda function and passing a SYCL item parameter. In this case both the global id and global range are queryable. This variant of parallel_for is designed for when it is necessary to query the global range within which the global id will vary. No information is queryable on the local (work-group) size chosen by the implementation.

```
1
   myQueue.submit([&](handler & cgh) {
2
        accessor acc { myBuffer, cgh, write_only };
3
4
        cgh.parallel_for(range<1>(numWorkItems),
5
                         [=] (item<1> item) {
            // kernel argument type is item
6
7
            size_t index = item.get_linear_id();
8
            acc[index] = index;
9
        });
10 });
```

Below is an example of invoking a SYCL kernel function with parallel_for using a lambda function and passing auto parameter, treated as item. In this case both the global id and global range are queryable. The same effect can be achieved using class with templatized operator(). This variant of parallel_for is designed for when it is necessary to query the global range within which the global id will vary. No information is queryable on the local (work-group) size chosen by the implementation.

```
myQueue.submit([&](handler & cgh) {
1
2
        auto acc = myBuffer.get_access<access::mode::write>(cgh);
3
4
        cgh.parallel_for(range<1>(numWorkItems),
5
                         [=] (auto item) {
6
            // kernel agrument type is auto treated as an item
7
            size_t index = item.get_linear_id();
8
            acc[index] = index;
9
        });
10 });
```

Below is an example of invoking a SYCL kernel function with parallel_for using a lambda function and passing integral type (e.g. int, size_t) parameter. This example is only valid when calling parallel_for with range<1>. In this case, only the global id is available. This variant of parallel_for is designed for when it is not necessary to query the global range of the index space being executed across, or the local (workgroup) size chosen by the implementation.

```
myQueue.submit([&](handler & cgh) {
1
2
       auto acc = myBuffer.get_access<access::mode::write>(cgh);
3
4
       cgh.parallel_for(range<1>(numWorkItems),
5
                        [=] (size_t index) {
6
           // kernel argument type is size_t
7
           acc[index] = index;
8
       });
9
  });
```

The parallel_for overload without an offset can be called with either a number or a braced-init-list with 1-3 elements. In that case the following calls are equivalent:

- parallel_for(N, some_kernel) has same effect as parallel_for(range<1>(N), some_kernel)
- parallel_for({N}, some_kernel) has same effect as parallel_for(range<1>(N), some_kernel)
- parallel_for({N1, N2}, some_kernel) has same effect as parallel_for(range<2>(N1, N2), some_kernel)
- parallel_for({N1, N2, N3}, some_kernel) has same effect as parallel_for(range<3>(N1, N2, N3), some_kernel)

Below is an example of invoking parallel_for with a number instead of an explicit range object.

```
1
   myQueue.submit([&](handler & cgh) {
2
        auto acc = myBuffer.get_access<access::mode::write>(cgh);
3
4
       // parallel_for may be called with number (with numWorkItems)
5
        cgh.parallel_for(numWorkItems,
6
                         [=] (auto item) {
7
            size_t index = item.get_linear_id();
8
            acc[index] = index;
9
       });
10 });
```

For SYCL kernel functions invoked via the above described overload of the parallel_for member function, it is disallowed to use local accessors or to use a work-group barrier.

The following two examples show how a kernel function object can be launched over a 3D grid, with 3 elements in each dimension. In the first case work-item ids range from 0 to 2 inclusive, and in the second case work-item ids run from 1 to 3.

```
1 myQueue.submit([&](handler & cgh) {
2
     cgh.parallel_for(
3
       range<3>(3,3,3), // global range
4
          [=] (item<3> it) {
5
            //[kernel code]
6
          });
7
   });
8
   myQueue.submit([&](handler & cgh) {
9
     cgh.parallel_for(
10
        range<3>(3,3,3), // global range
```

The last case of a parallel_for invocation enables low-level functionality of work-items and work-groups. This becomes valuable when an execution requires groups of work-items that communicate and synchronize. These are exposed in SYCL through parallel_for (nd_range,...) and the nd_item class. In this case, the developer needs to define the nd_range that the kernel will execute on in order to have fine grained control of the enqueuing of the kernel. This variation of parallel_for expects an nd_range, specifying both local and global ranges, defining the global number of work-items and the number in each cooperating work-group. The resulting function object is passed an nd_item instance making all the information available, as well as work-group barrier to synchronize between the work-items in the work-group.

The following example shows how sixty-four work-items may be launched in a three-dimensional grid with four in each dimension, and divided into eight work-groups. Each group of work-items synchronizes with a work-group barrier.

```
myQueue.submit([&](handler & cgh) {
1
2
     cgh.parallel_for(
3
         nd_range<3>(range<3>(4, 4, 4), range<3>(2, 2, 2)), [=](nd_item<3> item) {
4
           //[kernel code]
5
           // Internal synchronization
6
           group_barrier(item.get_group());
7
           //[kernel code]
8
         });
9
  });
```

In all of these cases the underlying nd-range will be created and the kernel defined as a function object will be created and enqueued as part of the command group scope.

A kernel_handler can optionally be passed as a parameter to the SYCL kernel function that is invoked by both variants of parallel_for.

```
myQueue.submit([&](handler & cgh) {
 1
 2
     cgh.parallel_for(
 3
        range<3>(3,3,3), // global range
 4
          [=] (item<3> it, kernel_handler kh) {
 5
            //[kernel code]
 6
          });
 7
    });
    myQueue.submit([&](handler & cgh) {
 8
 9
     cgh.parallel_for(
10
        range<3>(3,3,3), // global range
11
        id<3>(1,1,1), // offset
12
          [=] (item<3> it, kernel_handler kh) {
13
            //[kernel code]
14
          });
15 });
```

4.10.7.3 Parallel for hierarchical invoke

The hierarchical parallel kernel execution interface provides the same functionality as is available from the ndrange interface, but exposed differently. To execute the same sixty-four work-items in sixteen work-groups that we saw in the previous example, we execute an outer parallel_for_work_group call to create the groups. The member function handler::parallel_for_work_group is parameterized by the number of work-groups, such that the size of each group is chosen by the runtime, or by the number of work-groups and number of work-items for users who need more control.

The body of the outer **parallel_for_work_group** call consists of a lambda function or function object. The body of this function object contains code that is executed only once for the entire work-group. If the code has no side-effects and the compiler heuristic suggests that it is more efficient to do so, this code will be executed for each work-item.

Within this region any variable declared will have the semantics of local memory, shared between all work-items in the work-group. If the device compiler can prove that an array of such variables is accessed only by a single work-item throughout the lifetime of the work-group, for example if access is derived from the id of the work-item with no transformation, then it can allocate the data in private memory or registers instead.

To guarantee use of private per-work-item memory, the **private_memory** class can be used to wrap the data. This class simply constructs private data for a given group across the entire group. The id of the current work-item is passed to any access to grab the correct data.

The private_memory class has the following interface:

```
namespace sycl {
 1
    template <typename T, int Dimensions = 1>
2
3 class private_memory {
4
    public:
     // Construct based directly off the number of work-items
5
     private_memory(const group<Dimensions> &);
6
7
8
     // Access the instance for the current work-item
9
     T &operator()(const h_item<Dimensions> &id);
10 };
11 }
```

| Constructor | Description |
|---|--|
| <pre>private_memory(const group<dimensions> &)</dimensions></pre> | Place an object of type T in the underly- ing private memory of each work-items. The type T must be default constructible. The un- derlying constructor will be called for each work-item. |
| | End of table |

Table 4.94: Constructor of the private_memory class.

| Member functions | Description |
|---|--|
| <pre>T &operator()(const h_item<dimensions> &id)</dimensions></pre> | Retrieve a reference to the object for the work-items. |
| | End of table |

Table 4.95: Member functions of the private_memory class.

Private memory is allocated per underlying work-item, not per iteration of the parallel_for_work_item loop. The number of instances of a private memory object is only under direct control if a work-group size is passed to the parallel_for_work_group call. If the underlying work-group size is chosen by the runtime, the number of private memory instances is opaque to the program. Explicit private memory declarations should therefore be used with care and with a full understanding of which instances of a parallel_for_work_item loop will share the same underlying variable.

Also within the lambda body can be a sequence of calls to parallel_for_work_item. At the edges of these inner parallel executions the work-group synchronizes. As a result the pair of parallel_for_work_item calls in the code below is equivalent to the parallel execution with a work-group barrier in the earlier example.

```
myQueue.submit([&](handler & cgh) {
 1
 2
      // Issue 8 work-groups of 8 work-items each
 3
      cgh.parallel_for_work_group(
 4
          range<3>(2, 2, 2), range<3>(2, 2, 2), [=](group<3> myGroup) {
 5
 6
        //[workgroup code]
 7
        int myLocal; // this variable is shared between workitems
 8
        // this variable will be instantiated for each work-item separately
 9
        private_memory<int> myPrivate(myGroup);
10
11
        // Issue parallel work-items. The number issued per work-group is determined
12
        // by the work-group size range of parallel_for_work_group. In this case,
13
        // 8 work-items will execute the parallel_for_work_item body for each of the
14
        // 8 work-groups, resulting in 64 executions globally/total.
15
        myGroup.parallel_for_work_item([&](h_item<3> myItem) {
16
          //[work-item code]
17
          myPrivate(myItem) = 0;
18
        });
19
20
        // Implicit work-group barrier
21
22
        // Carry private value across loops
23
        myGroup.parallel_for_work_item([&](h_item<3> myItem) {
24
          //[work-item code]
25
          output[myItem.get_global_id()] = myPrivate(myItem);
26
        });
27
        //[workgroup code]
28
      });
29 });
```

It is valid to use more flexible dimensions of the work-item loops. In the following example we issue 8 workgroups but let the runtime choose their size, by not passing a work-group size to the parallel_for_work_group call. The parallel_for_work_item loops may also vary in size, with their execution ranges unrelated to the dimensions of the work-group, and the compiler generating an appropriate iteration space to fill the gap. In this case, the h_item provides access to local ids and ranges that reflect both kernel and parallel_for_work_item invocation ranges.

```
1
   myQueue.submit([&](handler & cgh) {
2
      // Issue 8 work-groups. The work-group size is chosen by the runtime because unspecified
3
      cgh.parallel_for_work_group(
 4
          range<3>(2, 2, 2), [=](group<3> myGroup) {
5
6
        // Launch a set of work-items for each work-group. The number of work-items is chosen
        /\!/ by the runtime because the work-group size was not specified to <code>parallel_for_work_group</code>
 7
8
        // and a logical range is not specified to parallel_for_work_item.
0
        myGroup.parallel_for_work_item([=](h_item<3> myItem) {
10
          //[work-item code]
11
        });
12
        // Implicit work-group barrier
13
14
15
        // Launch 512 logical work-items that will be executed by the underlying work-group size
16
        // chosen by the runtime. myItem allows the logical and physical work-item IDs to be
17
        // queried. 512 logical work-items will execute for each work-group, and the parallel_for
18
        // body will therefore be executed 8*512 = 4096 times globally/total.
19
        myGroup.parallel_for_work_item(range<3>(8, 8, 8), [=](h_item<3> myItem) {
          //[work-item code]
20
21
        });
22
        //[workgroup code]
23
      });
24 });
```

This interface offers a more intuitive way for tiling parallel programming paradigms. In summary, the hierarchical model allows a developer to distinguish the execution at work-group level and at work-item level using the parallel_for_work_group and the nested parallel_for_work_item functions. It also provides this visibility to the compiler without the need for difficult loop fission such that host execution may be more efficient.

A kernel_handler can optionally be passed as a parameter to the SYCL kernel function that is invoked by any variant of parallel_for_work_group.

```
1
   myQueue.submit([&](handler & cgh) {
2
      // Issue 8 work-groups of 8 work-items each
 3
      cgh.parallel_for_work_group(
 4
          range<3>(2, 2, 2), range<3>(2, 2, 2), [=](group<3> myGroup,
 5
          kernel_handler kh) {
6
7
        //[workgroup code]
8
        int myLocal; // this variable is shared between workitems
9
        // this variable will be instantiated for each work-item separately
10
        private_memory<int> myPrivate(myGroup);
11
12
        // Issue parallel work-items. The number issued per work-group is determined
13
        // by the work-group size range of parallel_for_work_group. In this case,
14
        // 8 work-items will execute the parallel_for_work_item body for each of the
15
        // 8 work-groups, resulting in 64 executions globally/total.
16
        myGroup.parallel_for_work_item([&](h_item<3> myItem) {
17
          //[work-item code]
```

```
18
          myPrivate(myItem) = 0;
19
        }):
20
21
        // Implicit work-group barrier
22
23
        // Carry private value across loops
24
        myGroup.parallel_for_work_item([&](h_item<3> myItem) {
25
          //[work-item code]
26
          output[myItem.get_global_id()] = myPrivate(myItem);
27
        });
28
        //[workgroup code]
29
      });
30 });
```

4.10.8 SYCL functions for explicit memory operations

In addition to kernels, command group objects can also be used to perform manual operations on host and device memory by using the *copy* API of the command group handler. Manual copy operations can be seen as specialized kernels executing on the device, except that typically this operations will be implemented using the OpenCL host API (e.g, enqueue copy operations).

The SYCL memory objects involved in a copy operation are specified using accessors. Explicit copy operations have a source and a destination. When an accessor is the *source* of the operation, the destination can be a host pointer or another accessor. The *source* accessor can have either read or read_write access mode.

When an accessor is the *destination* of the explicit copy operation, the source can be a host pointer or another accessor. The *destination* accessor can have either write, read_write, discard_write, discard_read_write access modes.

When accessors are both the origin and the destination, the operation is executed on objects controlled by the SYCL runtime. The SYCL runtime is allowed to not perfom an explicit in-copy operation if a different path to update the data is available according to the SYCL application memory model.

The most recent copy of the memory object may reside on any context controlled by the SYCL runtime, or on the host in a pointer controlled by the SYCL runtime. The SYCL runtime will ensure that data is copied to the destination once the command group has completed execution.

Whenever a host pointer is used as either the host or the destination of these explicit memory operations, it is the responsibility of the user for that pointer to have at least as much memory allocated as the accessor is giving access to, e.g: if an accessor accesses a range of 10 elements of int type, the host pointer must at least have 10 * sizeof(int) bytes of memory allocated.

A special case is the update_host member function. This member function only requires an accessor, and instructs the runtime to update the internal copy of the data in the host, if any. This is particularly useful when users use manual synchronization with host pointers, e.g. via mutex objects on the **buffer** constructors.

Table 4.96 describes the interface for the explicit copy operations.

| Member function | Description |
|--|--|
| <pre>template <typename access::mode<="" dim_src,="" int="" pre="" t_src,=""></typename></pre> | Copies the contents of the memory object |
| <pre>mode_src, access::target tgt_src, typename T_dest,</pre> | accessed by src into the memory pointed to |
| <pre>access::placeholder isPlaceholder></pre> | by dest. dest must have at least as many |
| <pre>void copy(accessor<t_src, dim_src,="" mode_src,<="" pre=""></t_src,></pre> | bytes as the range accessed by src. |
| <pre>tgt_src, isPlaceholder> src, std::shared_ptr<t_dest></t_dest></pre> | |
| dest) | |
| <pre>template <typename int<="" pre="" t_dest,="" t_src,="" typename=""></typename></pre> | Copies the contents of the memory pointed |
| <pre>dim_dest, access::mode mode_dest, access::target</pre> | to by src into the memory object accessed |
| <pre>tgt_dest, access::placeholder isPlaceholder></pre> | by dest. src must have at least as many |
| <pre>void copy(std::shared_ptr<t_src> src, accessor<</t_src></pre> | bytes as the range accessed by dest. |
| T_dest, dim_dest, mode_dest, tgt_dest, isPlaceholder | |
| > dest) | |
| <pre>template <typename access::mode<="" dim_src,="" int="" pre="" t_src,=""></typename></pre> | Copies the contents of the memory object |
| <pre>mode_src, access::target tgt_src, typename T_dest,</pre> | accessed by src into the memory pointed to |
| <pre>access::placeholder isPlaceholder></pre> | by dest. dest must have at least as many |
| <pre>void copy(accessor<t_src, dim_src,="" mode_src,<="" pre=""></t_src,></pre> | bytes as the range accessed by src. |
| <pre>tgt_src, isPlaceholder> src, T_dest * dest)</pre> | |
| <pre>template <typename int<="" pre="" t_dest,="" t_src,="" typename=""></typename></pre> | Copies the contents of the memory pointed |
| <pre>dim_dest, access::mode mode_dest, access::target</pre> | to by src into the memory object accessed |
| <pre>tgt_dest, access::placeholder isPlaceholder></pre> | by dest. src must have at least as many |
| <pre>void copy(const T_src * src, accessor<t_dest,< pre=""></t_dest,<></pre> | bytes as the range accessed by dest. |
| <pre>dim_dest, mode_dest, tgt_dest, isPlaceholder> dest)</pre> | |
| <pre>template <typename access::mode<="" dim_src,="" int="" pre="" t_src,=""></typename></pre> | Copies the contents of the memory object |
| <pre>mode_src, access::target tgt_src, access::</pre> | accessed by src into the memory object ac- |
| <pre>placeholder isPlaceholder_src, typename T_dest, int</pre> | cessed by dest. src must have at least as |
| <pre>dim_dest, access::mode mode_dest, access::target</pre> | many bytes as the range accessed by dest. |
| tgt_dest, access::placeholder isPlaceholder_dest> | |
| <pre>void copy(accessor<t_src, dim_src,="" mode_src,<="" pre=""></t_src,></pre> | |
| <pre>tgt_src, isPlaceholder_src> src, accessor<t_dest,< pre=""></t_dest,<></pre> | |
| <pre>dim_dest, mode_dest, tgt_dest, isPlaceholder_dest></pre> | |
| dest) | |
| template <typename access::mode="" dim,="" int="" mode,<="" t,="" td=""><td>The contents of the memory object accessed</td></typename> | The contents of the memory object accessed |
| access::target tgt, access::placeholder | via acc on the host are guaranteed to be up- |
| 1sPlaceholder> | to-date after this command group object ex- |
| void update_host(accessor<1, dim, mode, tgt, | ecution is complete. |
| 1splaceholder> acc) | |
| <pre>tempiate <typename 1,="" access::mode="" dim,="" int="" mode,<="" pre=""></typename></pre> | Replicates the value of src into the memory |
| access::target tgt, access::placenolder | object accessed by dest. I must be a scalar |
| 1SP1acenoloer> | value of a SYCL vector type. |
| void IIII(accessor<1, dlm, mode, tgt, | |
| isriacenoider> dest, | |
| | End of table |

Table 4.96: Member functions of the handler class.

The listing below illustrates how to use explicit copy operations in SYCL. The example copies half of the contents of a std::vector into the device, leaving the rest of the contents of the buffer on the device unchanged.

```
const size_t nElems = 10u;
 1
 2
3 // Create a vector and fill it with values 0, 1, 2, 3, 4, 5, 6, 7, 8, 9
 4 std::vector<int> v { nElems };
 5 std::iota(std::begin(v), std::end(v), 0);
 6
7
   // Create a buffer with no associated user storage
   sycl::buffer<int, 1> b { range<1>(nElems) };
8
9
10 // Create a queue
11 queue myQueue;
12
13 myQueue.submit([&](handler &cgh) {
14
     // Retrieve a ranged write accessor to a global buffer with access to the
      // first half of the buffer
15
16
     accessor acc { b, cgh, range<1>(nElems / 2), id<1>(0), write_only };
17
     // Copy the first five elements of the vector into the buffer associated with
18
     // the accessor
19
     cgh.copy(v.data(), acc);
20 });
```

4.10.9 Functions for using a module

```
1 void use_module(const module<module_state::executable> &execModule); // (1)
2
3 template <typename T>
4 void use_module(const module<module_state::executable> &execModule // (2)
5 T deviceImageSelector);
```

 Effects: The command group associated with the handler will use device images of the module execModule in any kernel invocation commands for all SYCL kernel functions represented by the module. If the module contains multiple device images that are compatible with the device associated with handler, then the device image chosen is implementation defined.

Throws: errc::invalid_object_error if the context associated with the command group handler via its associated queue is different from the context associated with the module specified by execModule.

2. *Effects:* The command group associated with the handler will use device images of the module execModule in any kernel invocation commands for all SYCL kernel functions represented by the module. All device images in the module which are compatible with the device associated with handler will be passed to the device image selection function. The device image with the highest score will be chosen.

Throws: errc::invalid_object_error if the context associated with the command group handler via its associated queue is different from the context associated with the module specified by execModule.

Throws: errc::device_image_selection_error if no device image could be selected.

4.10.10 Functions for using specialization constants

```
1 template<auto& S>
```

```
2 bool has_specialization_constant() const noexcept; // (1)
```

1. *Returns:* true if any of the SYCL kernel functions represented by the module associated with the command group handler contain the specialization constant represented by the specialization_id at the address S, otherwise returns false.

1 template<auto& S>

- 2 typename std::remove_reference_t<decltype(S)>::type get_specialization_constant(); // (1)
 - 1. *Returns:* From the module associated with the command group handler, the value of the specialization constant associated with the specialization_id at the address S if the specialization constant has been set, otherwise returns the default value.

Throws: errc::invalid_object_error if this->has_specialization_constant<S>() evaluates to false.

4.11 Host tasks

4.11.1 Overview

A host task is a native C++ callable which is scheduled by the SYCL runtime. A host task is submitted to a queue via a command group by a host task command.

When a host task command is submitted to a queue it is scheduled based on its data dependencies with other commands including kernel invocation commands and asynchronous copies, resolving any requisites created by accessors attached to the command group as defined in Section 3.7.1.

Since a host task is invoked directly by the SYCL runtime rather than being compiled as a SYCL kernel function, it does not have the same restrictions as a SYCL kernel function, and can therefore contain any arbitrary C++ code. However, capturing or using any SYCL class with reference semantics (see Section 4.5.3) is undefined behaviour.

A host task can be enqueued on any queue including a host queue and the callable will be invoked directly by the SYCL runtime, regardless of which device the queue is associated with.

A host task is enqueued on a queue via the host_task member function of the handler class.

A host task can optionally be used to interoperate with the native backend objects associated with the queue executing the host task, the context that the queue is associated with, the device that the queue is associated with and the accessors that have been captured in the callable, via an optional interop_handle parameter.

This allows host task to be used for two purposes: either as a task which can perform arbitrary C++ code within the scheduling of the SYCL runtime or as a task which can perform interoperability at a point within the scheduling of the SYCL runtime.

For the former use case host accessors should be used to request that a buffer or image be made available on the host so that it can be accessed directly via the accessor.

For the later use case device accessors should be used to request that a buffer or image be made available on the device associated with the queue used to submit the host task so that it can be accessed via interoperability member functions provided by the interop_handle class.

Local accessors cannot be used within a host task.

```
namespace sycl {
 1
 2
 3
    class interop_handle {
 4
     private:
 5
 6
      interop_handle(__unspecified__);
 7
 8
     public:
 9
10
      interop_handle() = delete;
11
12
      template <backend Backend, typename dataT, int dims, access::mode accessMode,
13
                access::target accessTarget, access::placeholder isPlaceholder>
14
      backend_traits<Backend>::native_type<buffer>
15
      get_native_mem(const accessor<dataT, dims, accessMode, accessTarget,</pre>
16
                                    isPlaceholder> &bufferAccessor) const;
17
18
      template <backend Backend, typename dataT, int dims, access::mode accessMode,
19
                access::target accessTarget, access::placeholder isPlaceholder>
20
      backend_traits<Backend>::native_type<image>
21
      get_native_mem(const accessor<dataT, dims, accessMode, accessTarget,</pre>
22
                                    isPlaceholder> &imageAccessor) const;
23
24
      template <backend Backend>
25
      backend_traits<Backend>::native_type<queue> get_native_queue() const noexcept;
26
27
      template <backend Backend>
28
      backend_traits<Backend>::native_type<device> get_native_device() const noexcept;
29
30
      template <backend Backend>
31
      backend_traits<Backend>::native_type<context> get_native_context() const noexcept;
32
33 };
34
35
    class handler {
36
      . . .
37
38
     public:
39
40
      template <typename T>
41
      void host_task(T &&hostTaskCallable);
42
43
      . . .
44 };
45
46 } // namespace sycl
```

4.11.2 Class interop_handle

The interop_handle class is an abstraction over the queue which is being used to invoke the host task and its associated device and context. It also represents the state of the SYCL runtime dependency model at the point the host task is invoked.

The interop_handle class provides access to the native backend object associated with the queue, device, context

and any buffers or images that are captured in the callable being invoked in order to allow a host task to be used for interoperability purposes.

An interop_handle cannot be constructed by user-code, only by the SYCL runtime.

1 class interop_handle;

4.11.2.1 Constructors

```
1 private:
2
3 interop_handle(__unspecified__); // (1)
4
5 public:
6
7 interop_handle() = delete; // (2)
```

- 1. Private implementation defined constructor with unspecified arguments so that the SYCL runtime can construct a interop_handle.
- 2. Explicitly deleted default constructor.

4.11.2.2 Template member functions get_native_*

```
1 template <backend Backend, typename dataT, int dims, access::mode accMode,</pre>
2
              access::target accTarget, access::placeholder isPlaceholder>
3 backend_traits<Backend>::native_type<buffer>
4 get_native_mem(const accessor<dataT, dims, accMode, accTarget, // (1)
5
                                  isPlaceholder> &bufferAccessor) const;
6
7
   template <backend Backend, typename dataT, int dims, access::mode accMode,
8
              access::target accTarget, access::placeholder isPlaceholder>
9
   backend_traits<Backend>::native_type<image>
   get_native_mem(const accessor<dataT, dims, accMode, accTarget, // (2)</pre>
10
11
                                  isPlaceholder> &imageAccessor) const;
12
13 template <backend Backend>
14 backend_traits<Backend>::native_type<queue> get_native_queue() const noexcept; // (3)
15
16 template <backend Backend>
17 backend_traits<Backend>::native_type<device> get_native_device() const noexcept; // (4)
18
19 template <backend Backend>
20 backend_traits<Backend>::native_type<context> get_native_context() const noexcept; // (5)
```

- - 1. *Constraints:* Available only if the optional interoperability function get_native taking a buffer is available and if accTarget is access::target::global_buffer or access::target::constant_buffer.

Returns: The SYCL application interoperability native backend object associated with the accessor bufferAccessor. The native backend object returned must be in a state where it represents the memory in its current state within the SYCL runtime dependency model be capable of being used in a way appro-

priate for the associated SYCL backend. It is undefined behaviour to use the native backend object outside of the scope of the host task.

Throws: errc::invalid_object_error if the accessor bufferAccessor was not registered with the command group which contained the host task.

 Constraints: Available only if the optional interoperability function get_native taking an unsampled_image or sampled_image is available and if accTarget is access::target::unsampled_image or access::target ::sampled_image.

Returns: The SYCL application interoperability native backend object associated with the accessor imageAccessor. The native backend object returned must be in a state where it represents the memory in its current state within the SYCL runtime dependency model and is capable of being used in a way appropriate for the associated SYCL backend. It is undefined behaviour to use the native backend object outside of the scope of the host task.

Throws: errc::invalid_object_error if the accessor imageAccessor was not registered with the command group which contained the host task.

3. Constraints: Available only if the optional interoperability function get_native taking a queue is available.

Returns: The SYCL application interoperability native backend object associated with the queue that the host task was submitted to. If the command group was submitted with a secondary queue and the fall-back was triggered, the queue that is associated with the interop_handle must be the fall-back queue. The native backend object returned must be in a state where it is capable of being used in a way appropriate for the associated SYCL backend. It is undefined behaviour to use the native backend object outside of the scope of the host task.

4. *Constraints:* Available only if the optional interoperability function get_native taking a device is available.

Returns: The SYCL application interoperability native backend object associated with the device that is associated with the queue that the host task was submitted to. The native backend object returned must be in a state where it is capable of being used in a way appropriate for the associated SYCL backend. It is undefined behaviour to use the native backend object outside of the scope of the host task.

5. *Constraints:* Available only if the optional interoperability function get_native taking a context is available.

Returns: The SYCL application interoperability native backend object associated with the context that is associated with the queue that the host task was submitted to. The native backend object returned must be in a state where it is capable of being used in a way appropriate for the associated SYCL backend. It is undefined behaviour to use the native backend object outside of the scope of the host task.

4.11.3 Additions to the handler class

This section describes member functions in the command group handler class that are used with host tasks.

```
1 class handler {
2 ...
3
4 public:
```

CHAPTER 4. SYCL PROGRAMMING INTERFACE

266

```
5 template <typename T>
6 void host_task(T &&hostTaskCallable); // (1)
7
8 ...
9 };
```

1. *Effects:* Enqueues an implementation defined command to the SYCL runtime to invoke hostTaskCallable exactly once. The scheduling of the invocation of hostTaskCallable in relation to other commands enqueued to the SYCL runtime must be in accordance with the dependency model described in Section 3.7.1. Initialises an interop_handle object and passes it to hostTaskCallable when it is invoked if std ::is_invocable_v<T, interop_handle> evaluates to true, otherwise invokes invokes hostTaskCallable as a nullary function.

4.12 Kernel class

The kernel class is an abstraction of a kernel object in SYCL. In the most common case the kernel object will contain the compiled version of a kernel invoked inside a command group using one of the parallel interface functions as described in 4.10.7. The SYCL runtime will create a kernel object, when it needs to enqueue the kernel on a command queue.

In the case where a developer would like to pre-compile a kernel or compile and link it with an existing program, then the kernel object will be created and contain that kernel using the module class, as defined in 4.13.6. In both of the above cases, the developer cannot instantiate a kernel object but can instantiate a named function object type that they could use, or create a function object from a kernel member function using C++ features. The kernel class object needs a parallel_for(...) invocation or an explicitly built SYCL kernel instance, for this compilation of the kernel to be triggered.

The SYCL kernel class provides the common reference semantics (see Section 4.5.3).

The member functions of the SYCL kernel class are listed in Table 4.97. The additional common special member functions and common member functions are listed in Tables 4.1 and 4.2, respectively.

```
1
   namespace sycl {
    class kernel {
2
3
     private:
 4
      kernel();
 5
6
     public:
7
      /* -- common interface members -- */
8
9
10
      backend get_backend() const;
11
12
      bool is_host() const;
13
14
      context get_context() const;
15
16
      module<module_state::executable> get_module() const;
17
18
      template <info::kernel param>
19
      typename info::param_traits<info::kernel, param>::return_type
20
      get_info() const;
```

| 21 22 23 24 25 | <pre>template <info::kernel_device_specific param=""> typename info::param_traits<info::kernel_device_specific, param="">::return_type get_info(const device &dev) const;</info::kernel_device_specific,></info::kernel_device_specific></pre> |
|----------------------------|--|
| 26 | template <typename backendenum="" backendenum,="" param=""></typename> |
| 27 | <pre>typename info::param_traits<backendenum, param="">::return_type</backendenum,></pre> |
| 28 | <pre>get_backend_info() const;</pre> |
| 29 | |
| 30 | <pre>template <info::kernel_work_group param=""></info::kernel_work_group></pre> |
| 31 | <pre>typename info::param_traits<info::kernel_work_group, param="">::return_type</info::kernel_work_group,></pre> |
| 32 | <pre>get_work_group_info(const device &dev) const;</pre> |
| 33 | }; |
| 34 | } // namespace sycl |

| Member functions | Description |
|--|---|
| <pre>backend get_backend()const</pre> | Returns the backend identifying the SYCL |
| | backend associated with this kernel. |
| <pre>bool is_host()const</pre> | Returns true if this SYCL kernel is a host |
| | kernel. |
| <pre>context get_context()const</pre> | Return the context that this kernel is de- |
| | fined for. The value returned must be equal |
| | to that returned by get_info <info::kernel< td=""></info::kernel<> |
| | ::context>(). |
| <pre>module<module_state::executable> get_module()const</module_state::executable></pre> | Returns the module that this kernel is part of. |
| | The value returned must be equal to that re- |
| | <pre>turned by get_info<info::kernel::module< pre=""></info::kernel::module<></pre> |
| | >(). |
| <pre>template <info::kernel param=""></info::kernel></pre> | Query information from the kernel object |
| <pre>typename info::param_traits<</pre> | using the info::kernel descriptor. |
| <pre>info::kernel, param>::return_type</pre> | |
| <pre>get_info()const</pre> | |
| <pre>template <info::kernel_device_specific param=""></info::kernel_device_specific></pre> | Query information from a kernel using the |
| <pre>typename info::param_traits<</pre> | <pre>info::kernel_device_specific descriptor</pre> |
| <pre>info::kernel_device_specific, param>::</pre> | for a specific device. |
| return_type | |
| <pre>get_info(const device &dev)const</pre> | |
| | Continued on next page |

Table 4.97: Member functions of the kernel class.

| Member functions | Description |
|---|--|
| <pre>template <typename backendenum="" backendenum,="" param=""></typename></pre> | Queries this SYCL kernel for SYCL back- |
| <pre>typename info::param_traits<backendenum, param="">::</backendenum,></pre> | end-specific information requested by the |
| return_type | template parameter param. BackendEnum |
| <pre>get_backend_info()const</pre> | can be any enum class type specified |
| | by the SYCL backend specification of a |
| | supported SYCL backend named accord- |
| | ing to the convention info:: <backend_name< td=""></backend_name<> |
| | >::kernel and param must be a valid |
| | enumeration of that enum class. Spe- |
| | cializations of info::param_traits must |
| | be defined for BackendEnum in accor- |
| | dance with the SYCL backend specifica- |
| | tion. Must throw an exception with the |
| | <pre>errc::invalid_object_error error code if</pre> |
| | the SYCL backend that corresponds with |
| | BackendEnum is different from the SYCL |
| | backend that is associated with this kernel |
| | • |
| <pre>template <info::kernel_work_group param=""></info::kernel_work_group></pre> | Query information from the work- |
| <pre>typename info::param_traits<</pre> | group from a kernel using the info:: |
| info::kernel_work_group, param>::return_type | kernel_work_group descriptor for a specific |
| <pre>get_work_group_info(const device &dev)const</pre> | device. This query has been deprecated in |
| | SYCL 2020 and will likely be removed in a |
| | future version of SYCL. |
| | End of table |

Table 4.97: Member functions of the kernel class.

4.12.1 Kernel information descriptors

A kernel can be queried for information using the get_info member function of the kernel class, specifying one of the info parameters enumerated in info::kernel. Every kernel (including a host kernel) must produce a valid value for each info parameter. The possible values for each info parameter and any restriction are defined in the specification of the SYCL backend associated with the kernel. All info parameters in info::kernel are specified in Table 4.98 and the synopsis for info::kernel is described in appendix A.5.

| Kernel Descriptors | Return type | Description |
|--|---------------------------|--|
| <pre>info::kernel::function_name</pre> | <pre>std::string</pre> | Return the kernel function name. |
| <pre>info::kernel::num_args</pre> | uint32_t | Return the number of arguments of the ex- |
| | | tracted kernel. |
| <pre>info::kernel::context</pre> | context | Return the SYCL context associated with this |
| | | SYCL kernel. |
| <pre>info::kernel::module</pre> | module< | Return the SYCL module associated with this |
| | <pre>module_state::</pre> | SYCL kernel. |
| | <pre>executable></pre> | |
| <pre>info::kernel::attributes</pre> | <pre>std::string</pre> | Return any attributes specified on a kernel |
| | | function (as defined in Section 5.7). |
| | | End of table |

Table 4.98: Kernel class information descriptors.

A kernel can also be queried for device-specific information using the get_info member function of the kernel class, specifying one of the info parameters enumerated in info::kernel_device_specific. Every kernel (including a host kernel) must produce a valid value for each info parameter. The possible values for each info parameter and any restriction are defined in the specification of the SYCL backend associated with the kernel. All info parameters in info::kernel_device_specific are specified in Table 4.99. The synopsis for info::kernel_device_specific is described in appendix A.5.

| Device-specific Kernel Information | Return type | Description |
|---|-------------|---|
| Descriptors | | |
| <pre>info::kernel_device_specific::</pre> | range<3> | Returns the maximum global work size. Only |
| global_work_size | | valid if device is of device_type custom or the |
| | | kernel is a built-in kernel. |
| <pre>info::kernel_device_specific::</pre> | size_t | Returns the maximum work-group size that can |
| work_group_size | | be used to execute a kernel on a specific device. |
| <pre>info::kernel_device_specific::</pre> | range<3> | Returns the work-group size specified by the |
| <pre>compile_work_group_size</pre> | | device compiler if applicable, otherwise returns |
| | | (0, 0, 0) |
| <pre>info::kernel_device_specific::</pre> | size_t | Returns a value, of which work-group size |
| <pre>preferred_work_group_size_multiple</pre> | | is preferred to be a multiple, for execut- |
| | | ing a kernel on a particular device. This |
| | | is a performance hint. The value must be |
| | | less than or equal to that returned by info:: |
| | | <pre>kernel_device_specific::work_group_size.</pre> |
| <pre>info::kernel_device_specific::</pre> | size_t | Returns the minimum amount of private mem- |
| <pre>private_mem_size</pre> | | ory, in bytes, used by each work-item in the |
| | | kernel. This value may include any private |
| | | memory needed by an implementation to ex- |
| | | ecute the kernel, including that used by the lan- |
| | | guage built-ins and variables declared inside |
| | | the kernel in the private address space. |
| <pre>info::kernel_device_specific::</pre> | uint32_t | Returns the maximum number of sub-groups |
| <pre>max_num_sub_groups</pre> | | for this kernel. |
| <pre>info::kernel_device_specific::</pre> | uint32_t | Returns the number of sub-groups specified by |
| compile_num_sub_groups | | the kernel, or 0 (if not specified). |
| <pre>info::kernel_device_specific::</pre> | uint32_t | Returns the maximum sub-group size for this |
| <pre>max_sub_group_size</pre> | | kernel. |
| <pre>info::kernel_device_specific::</pre> | uint32_t | Returns the required sub-group size specified |
| <pre>compile_sub_group_size</pre> | | by the kernel, or 0 (if not specified). |
| | | End of table |

Table 4.99: Device-specific kernel information descriptors.

Alternatively, a kernel can be queried for work-group information using the get_work_group_info member functions of the kernel class, specifying one of the info parameters enumerated in info::kernel_work_group. This query has been deprecated in SYCL 2020, and will likely be removed in a future version of SYCL.

| Kernel Work-group Information De- | Return type | Description |
|---|-------------|---|
| scriptors | | |
| <pre>info::kernel_work_group::</pre> | range<3> | Returns the maximum global work size. Only |
| global_work_size | | valid if device is of device_type custom or the |
| | | kernel is a built-in kernel. |
| <pre>info::kernel_work_group::</pre> | size_t | Returns the maximum work-group size that can |
| work_group_size | | be used to execute a kernel on a specific device. |
| <pre>info::kernel_work_group::</pre> | range<3> | Returns the work-group size specified by the |
| compile_work_group_size | | device compiler if applicable, otherwise returns |
| | | (0, 0, 0) |
| <pre>info::kernel_work_group::</pre> | size_t | Returns a value, of which work-group size |
| <pre>preferred_work_group_size_multiple</pre> | | is preferred to be a multiple, for execut- |
| | | ing a kernel on a particular device. This |
| | | is a performance hint. The value must be |
| | | less than or equal to that returned by info:: |
| | | <pre>kernel_work_group::work_group_size.</pre> |
| <pre>info::kernel_work_group::</pre> | size_t | Returns the minimum amount of private mem- |
| private_mem_size | | ory, in bytes, used by each work-item in the |
| | | kernel. This value may include any private |
| | | memory needed by an implementation to ex- |
| | | ecute the kernel, including that used by the lan- |
| | | guage built-ins and variables declared inside |
| | | the kernel in the private address space. |
| | | End of table |

Table 4.100: Kernel work-group information descriptors.

4.13 Modules

4.13.1 Overview

A module is a high-level abstraction which represents a set of SYCL kernel functions which are associated with a context and can be executed on a number of devices, where each device is associated with that same context.

The SYCL kernel functions represented by a module can exist within one or more device images of implementation defined file formats. Each device image within a module must contain the necessary symbols and meta-data for each SYCL kernel function that the containing module represents.

[Note: For example, a module associated with a context of an OpenCL SYCL backend, that represents the SYCL kernel function foo could contain two modules; one of SPIR-V and one of a vendor specific ISA, both containing foo in the relevant file format. — end note]

Modules are used to express the explicit compilation of SYCL kernel functions in order to then be invoked via a kernel invocation command such as parallel_for. Normally this compilation is done implicitly by the kernel invocation command, however it can be useful to perform the compilation manually in order to add custom properties to the compilation or to link SYCL kernel functions with other libraries.

A module can be obtained either by requesting the module associated with the current translation unit or via some SYCL backend-specific operation.

Once a module has been obtained there are a number of free functions for performing compilation, linking and joining.

A module can then be bound to a command group so that the SYCL kernel functions represented by the module are used in any kernel invocation commands. See Section 4.10.9 for more details.

4.13.2 Specialization constants

A SYCL kernel function can contain specialization constants which represent a constant variable where the value is not known until compilation of the SYCL kernel function. Any specialization constants of a given SYCL kernel function are exposed via an input module representing that SYCL kernel function, where the value of the specialization constants can be set. However, once a module has been compiled, resulting in an object module or executable module the value of specialization constants can no longer be changed.

Any specialization constants in a SYCL kernel function which are not set before that function is invoked will take a default value. This includes invoking a kernel invocation command such as parallel_for or retrieving an object module or executable module directly triggering implicit compilation.

As a module may contain more than one device image, some of these may natively support specialization constants and some may not, however all device images must set the value of specialization constants.

[Note: It is expected that a specialization constant is implemented either via an implementation-defined mechanism available to the file format of that device image such as SPIR-V specialization constants or by passing the value as an additional kernel argument to the SYCL kernel function via the SYCL runtime when it's invoked. — end note]

A specialization id is an identifier which represents a reference to a specialization constant both in the SYCL application for setting the value prior to the compilation of an input module and in a SYCL kernel function for retrieving the value during invocation.

A module that is associated with a host context may not contain native-specialization constants, though it must still emulate all specialization constants the SYCL kernel functions contains.

4.13.3 Synopsis

```
1
    namespace sycl {
 2
 3
    template <typename T>
 4
    class specialization_id {
 5
     private:
 6
 7
      specialization_id(const specialization_id& rhs) = delete;
 8
 9
      specialization_id(specialization_id&& rhs) = delete;
10
      specialization_id & operator=(const specialization_id& rhs) = delete;
11
12
13
      specialization_id &operator=(specialization_id&& rhs) = delete;
14
15
     public:
16
17
      using value_type = T;
```

```
18
19
      template<class... Args >
20
      explicit constexpr specialization_id(Args&&... args);
21 };
22
23 enum class module_state {
24
     input,
25
      object.
26
     executable
27 };
28
29
   template<module_state State>
30 class module {
31
    private:
32
33
     module(__unspecified__);
34
35
     public:
36
37
      using device_image_type = __unspecified__;
38
39
      using device_image_iterator = __unspecified__;
40
41
      module() = delete;
42
43
      context get_context() const noexcept;
44
45
      std::vector<device> get_devices() const noexcept;
46
47
      bool has_kernel(std::string kernelName) const noexcept;
48
49
      kernel get_kernel(std::string kernelName) const;
50
51
      std::vector<std::string> get_kernel_names() const;
52
53
      bool is_empty() const noexcept;
54
55
      device_image_iterator begin() const;
56
57
      device_image_iterator end() const;
58
59
      bool contains_specialization_constants() const noexcept;
60
61
      bool native_specialization_constant() const noexcept;
62
63
      template<auto& S>
64
      bool has_specialization_constant() const noexcept;
65
66
      template<auto& S>
67
      void set_specialization_constant(
68
        typename std::remove_reference_t<decltype(S)>::type value);
69
70
      template<auto& S>
71
      typename std::remove_reference_t<decltype(S)>::type get_specialization_constant() const;
72 };
```

```
73
74
    module<module state::executable>
75
    build(const module<module_state::input> &inputModule,
76
           const property_list &propList = {});
77
78
    module<module_state::object>
79
    compile(const module<module_state::input> &inputModule,
80
            const property_list &propList = {});
81
82 module<module_state::executable>
83
    link(const module<module_state::object> &objModule,
         const property_list &propList = {});
84
85
86
    module<module_state::executable>
    link(const std::vector<module<module_state::object>> &objModules,
87
88
         const property_list &propList = {});
89
90 template<module_state T>
91
    module<T> join(const std::vector<module<T>> &modules);
92
93 namespace this_module {
94
95 template <typename T>
96
    std::string kernel_name_v;
97
98
    bool has_any_module(context ctxt);
99
100 template<module_state S>
101
    bool has_module_in(context ctxt);
102
103 template<module_state S>
104
    module<T> get(context ctxt);
105
106 } // this_module
107
108 } // namespace sycl
```

4.13.4 Enum class module_state

A module can be in one of three different module states; input, object and executable. The module states reflect the current state of the device images and subsequently the SYCL kernel functions. The module states also alters the capabilities of the module.

The three module states are represented by an enum class called module_state.

```
1 enum class module_state {
2    input,
3    object,
4    executable
5 };
```

The values of each enumeration of module_state are implementation defined.

4.13.5 Class template specialization_id

A specialization id is represented by the class template specialization_id with a single template parameter T which specifies the unique name used to identify the associated specialization constant.

The template parameter T must be a forward-declarable type and specialization_id objects must be declared with automatic or static storage duration within a namespace or class scope.

```
1 template <typename T>
```

```
2 class specialization_id;
```

4.13.5.1 Constructors

```
1 template<class... Args>
```

- 2 explicit constexpr specialization_id(Args&&... args);
 - 1. *Constraints:* Available only when std::is_constructible_v<T, Args...> evaluates to true.

Effects: Constructs a specialization_id containing an instance of T initialized with args..., which represents the default value of the specialization constant.

4.13.5.2 Special member functions

```
1 specialization_id(const specialization_id& rhs) = delete; // (1)
2
3 specialization_id(specialization_id&& rhs) = delete; // (2)
4
5 specialization_id &operator=(const specialization_id& rhs) = delete; // (3)
6
7 specialization_id &operator=(specialization_id&& rhs) = delete; // (4)
```

- 1. Deleted copy constructor.
- 2. Deleted move constructor.
- 3. Deleted copy assignment operator.
- 4. Deleted move assignment operator.

4.13.6 Class template module

A module is represented by the class template module with the single template parameter State of type module_state which specifies its module state.

There are three different module types, to reflect the three module states:

- An input module is a module of the input module state and represents SYCL kernel functions which are yet to be compiled such as a source or intermediate representation.
- A object module is a module of the object module state and represents SYCL kernel functions which have been compiled but are yet to be linked such as an intermediate object of the ISAs of the associated devices.

• A executable module is a module of the executable module state and represents SYCL kernel functions which have been compiled and linked such as an executable of the ISAs of the associated devices.

A module cannot be constructed by user-code, only by the SYCL runtime.

A module is permitted to be empty in which case it contains no device images and represents no SYCL kernel functions.

A module is considered to have reference semantics as specified in Section 4.5.3, therefore any module constructed as a copy of another and the module that was copied from are considered to be equal. Furthermore, two modules of the same module state are considered to be equal if they are associated with the same context and devices and contain the same device images and subsequently the same SYCL kernel function.

A module must contain a copy of the context and devices that are associated with it for the duration of its lifetime. This means that the destructor of the associated context or devices will not be invoked if the module is still alive in accordance with Section 4.5.3.

```
1 template<module_state State>
2 class module;
```

4.13.6.1 Constructors

```
1 private:
2
3 module(__unspecified__); // (1)
4
5 public:
6
7 module() = delete; // (2)
```

- 1. Private implementation defined constructor with unspecified arguments so that the SYCL runtime can construct a module.
- 2. Explicitly deleted default constructor.

4.13.6.2 Member functions

- 1 context get_context() const noexcept; // (1)
 - 1. Returns: A context object representing the associated context.
- 1 std::vector<device> get_devices() const noexcept; // (1)
 - 1. Returns: A std::vector of device objects representing the associated devices.
- 1 bool has_kernel(std::string kernelName) const noexcept; // (1)
 - 1. *Constraints:* Available only when State == module_state::executable.

CHAPTER 4. SYCL PROGRAMMING INTERFACE

276

Returns: true if the module represents a SYCL kernel function with the value of string kernel name kernelName, otherwise returns false. Only available when the module is a executable module.

- 1 kernel get_kernel(std::string kernelName) const; // (1)
 - 1. *Constraints:* Available only when State == module_state::executable.

Returns: a kernel object representing the SYCL kernel function with the string kernel name with the value of kernelName if this->has_kernel(kernelName) evaluates to true, otherwise throws exception with the errc::invalid_object error code.

- 1 std::vector<std::string> get_kernel_names() const; // (1)
 - 1. *Returns:* A std::vector of std::string objects representing each of the SYCL kernel functions represented by the module.
- 1 bool is_empty() const noexcept; // (1)
 - 1. Returns: true if the module contains no device images, otherwise returns false.
- 1 device_image_iterator begin() const; // (1)
- 3 device_image_iterator end() const; // (2)
 - 1. *Returns:* An iterator of type device_image_iterator pointing to the beginning of a sequence of device images of type std::iterator_traits<device_image_iterator>::value_type.
 - 2. *Returns:* An iterator of type device_image_iterator pointing to the end of a sequence of device images of type std::iterator_traits<device_image_iterator>::value_type.
- 1 bool contains_specialization_constants() const noexcept; // (1)
 - 1. *Returns:* true if any SYCL kernel function represented by the module contains a specialization constant, otherwise returns false.
- 1 bool native_specialization_constant() const noexcept; // (1)
 - 1. *Returns:* true if all of the specialization constants contained in the module support are native-specialization constants for all device images.

2

- 2 bool has_specialization_constant() const noexcept; // (1)
 - 1. *Returns:* true if any of the SYCL kernel functions represented by the module contains the specialization constant represented by the specialization_id at the address S, otherwise returns false.

¹ template<auto& S>

- 1 template<auto& S>
- 2 void set_specialization_constant(
- 3 typename std::remove_reference_t<decltype(S)>::type value); // (1)
 - 1. *Constraints:* Available only when State == module_state::input.

Effects: Sets the value of the specialization constant represented by the specialization_id at the address S in all SYCL kernel functions which contain that specialization constant and for all device images. If the specialization constant was already set, then the previous value is overwritten. If two or more SYCL kernel functions contain the same specialization constant they are assumed to be the same and will have the same value.

Throws: errc::invalid_object_error if this->has_specialization_constant<S>() evaluates to false.

1 template<auto& S>

- 2 typename std::remove_reference_t<decltype(S)>::type get_specialization_constant() const; // (1)
 - 1. *Returns:* the value of the specialization constant associated with the specialization_id at the address S if the specialization constant has been set, otherwise returns the default value.

Throws: errc::invalid_object_error if this->has_specialization_constant<S>() evaluates to false.

4.13.7 Free functions

Modules can be compiled, linked, built or joined together using free functions which operate on modules of different module states.

When a module is being created as a results of one of these operations it is permitted to perform ad-hoc implementation-defined operations such as just-in-time compilation or translations to alter the file format of the device image in order to create a module of a resulting module state, however this is not required.

[Note: For example, if a SYCL kernel function was compiled by a device compiler to generate the file format SPIR-V, a module associated with a context of an OpenCL SYCL backend could be created as an input module using the SPIR-V file format directly or it could be created as a executable module implicitly triggering online compilation via the OpenCL runtime. — end note]

```
1 namespace sycl {
 2
 3 module<module_state::executable>
 4
    build(const module<module_state::input> &inputModule,
 5
          const property_list &propList = {}); // (1)
 6
 7
    module<module_state::object>
 8
    compile(const module<module_state::input> &inputModule,
 9
            const property_list &propList = {}); // (2)
10
11 module<module_state::executable>
    link(const module<module_state::object> &objModule,
12
13
         const property_list &propList = {}); // (3)
14
15 module<module_state::executable>
16 link(const std::vector<module<module_state::object>> &objModules,
```

```
17 const property_list &propList = {}); // (4)
18
19 template<module_state T>
20 module<T> join(const std::vector<module<T>> &modules); // (5)
21
22 } // namespace sycl
```

1. *Effects:* Performs implementation defined build operation(s), including compilation and linking, on the input module input, applying any properties provided via propList.

Returns: A module_state::executable> object containing the result of the build operation(s) performed. The module returned must represent the same SYCL kernel functions and be associated with the same context and devices as that of inputModule, however the device images may differ.

Throws: errc::build_error if none of the devices associated with the module have aspect:: online_compiler.

Throws: errc::build_error if the build operation(s) fail.

2. *Effects:* Performs implementation defined compilation operation(s) on the input module inputModule, applying any properties provided via propList.

Returns: A module_state::object> object containing the result of the compilation operation(s) performed. The module returned must represent the same SYCL kernel functions and be associated with the same context and devices as that of inputModule, however the device images may differ.

Throws: errc::compile_error if none of the devices associated with the module have aspect:: online_compiler.

Throws: errc::compile_error if the compilation operation(s) fail.

3. *Effects:* Performs implementation defined linking operation(s) on the input module objModule, applying any properties provided via propList. If two or more modules contain the same SYCL kernel functions it is assumed that they are the same so one of them is selected and the rest are discarded. The one which is selected is implementation defined.

Returns: A module_state::executable> object containing the result of the linking operation(s) performed. The module returned must represent the same SYCL kernel functions and be associated with the same context and devices as that of objModule, however the device images may differ.

Throws: errc::link_error if none of the devices associated with the module have aspect:: online_linker.

Throws: errc::link_error if the compilation operation(s) fail.

4. *Preconditions:* Each module in moduleObjects must be associated with the same context and devices.

Effects: Performs implementation defined linking operation(s) on the input modules in moduleObjects, applying any properties provided via propList. If two or more modules contain the same SYCL kernel functions it is assumed that they are the same so one of them is selected and the rest are discarded. The one which is selected is implementation defined.

Returns: A module_state::executable> object containing the result of the linking operation(s)

performed. The module returned must represent the same SYCL kernel functions and be associated with the same context and devices as that of objModules, however the device images may differ.

Throws: errc::link_error if none of the devices associated with the module have aspect:: online_linker.

Throws: errc::link_error if the compilation operation(s) fail.

Throws: errc::invalid_object_error if objModules.empty() evaluates to true or any module in objModules is associated with a different context or devices than another module in objModules.

5. Preconditions: Each module in modules must be associated with the same context and devices.

Effects: Performs implementation defined joining operation(s) on the modules in modules, applying any properties provided via propList. If two or more modules contain the same SYCL kernel functions it is assumed that they are the same so one of them is selected and the rest are discarded. The one which is selected is implementation defined.

Returns: A module object of module_state T, containing the result of the linking operation(s) performed. The module returned must represent all of the combined SYCL kernel functions associated with each module in modules and must be associated with the same context and devices as that of modules, however the device images may differ.

Throws: errc::link_error if the joining operation(s) fail.

Throws: errc::invalid_object_error if modules.empty() evaluates to true or any module in modules is associated with a different context or devices than another module in modules.

4.13.8 Namespace this_module

The namespace this_module provides additional free functions which can be used for querying and retrieving modules available to the current translation unit.

4.13.8.1 Type traits

```
1 template <typename T>
2 std::string kernel_name_v; // (1)
```

1. Template variable that takes a type T specifying the type kernel name of a SYCL kernel function in the current translation unit and whose value is the corresponding string kernel name.

4.13.8.2 Free functions

```
1 bool has_any_module(context ctxt); // (1)
2
3 template<module_state S> // (2)
4 bool has_module_in(context ctxt);
5
6 template<module_state S> // (3)
7 module<T> get(context ctxt);
```

- 1. *Returns:* true if there is a module of any module state, available within the current translation unit that is compatible with the context ctxt, otherwise returns false. Must not perform any compilation, linking, building or joining operation(s).
- 2. *Returns:* true if there is a module of the module state specified by the module_state value S, available or retrievable within the current translation unit that is compatible with the context ctxt, otherwise returns false. Must not perform any compilation, linking, building or joining operation(s).
- 3. *Effects:* May perform implementation defined compilation, linking or building operations in order to transform a module that is available in another module state into the module state specified by the module_state value S

Returns: A module of module_state S representing a module associated with the current translation unit and the context ctxt if this_module::has_module_in<S>() evaluates to true, otherwise returns a module that is empty. If the returned module is not empty, it must represent the set of SYCL kernel functions available to the current translation unit and may also contain a further implementation defined set of SYCL kernel functions.

[Note: A module returned from module<State>::get may contain additional SYCL kernel functions to those available in the current translation unit in order to facilitate different single-source compilation methods. This does not impact user code as link and join are required to assume any duplicate SYCL kernel functions are the same. — end note]

4.14 Defining kernels

In SYCL, functions that are executed on a SYCL device are referred to as SYCL kernel functions. A kernel containing such a SYCL kernel function is enqueued on a device queue in order to be executed on that particular device.

The return type of the SYCL kernel function is void, and all memory accesses between host and device are through the accessor class (Section 4.7.6) or through USM pointers (Section 4.8).

There are three ways of defining kernels: as named function objects, as lambda functions, or through backendspecific interoperability interfaces for modules and kernels, where available.

4.14.1 Defining kernels as named function objects

A kernel can be defined as a named function object type. These function objects provide the same functionality as any C++ function object, with the restriction that they need to follow C++ rules to be trivially copyable. The kernel function can be templated via templating the kernel function object type. For details on restrictions for kernel naming, please refer to Section 5.2.

The operator() member function must be const-qualified, and it may take different parameters depending on the data accesses defined for the specific kernel. If the operator() function writes to any of the member variables, the behavior is undefined.

The following example defines a SYCL kernel function, *RandomFiller*, which initializes a buffer with a random number. The random number is generated during the construction of the function object while processing the command group. The operator() member function of the function object receives an item object. This member function will be called for each work item of the execution range. The value of the random number will be assigned to each element of the buffer. In this case, the accessor and the scalar random number are members of

the function object and therefore will be parameters to the device kernel. Usual restrictions of passing parameters to kernels apply.

```
class RandomFiller {
 1
 2
     public:
 3
      RandomFiller(accessor<int> ptr)
 4
          : ptr_ { ptr } {
 5
        std::random_device hwRand;
 6
        std::uniform_int_distribution<> r { 1, 100 };
 7
        randomNum_ = r(hwRand);
 8
      }
 9
      void operator()(item<1> item) const { ptr_[item.get_id()] = get_random(); }
10
      int get_random() { return randomNum_; }
11
12
     private:
      accessor<int> ptr_;
13
14
      int randomNum_;
15
    };
16
17
    void workFunction(buffer<int, 1>& b, queue& q, const range<1> r) {
18
        myQueue.submit([&](handler& cgh) {
19
          accessor ptr { buf, cgh };
20
          RandomFiller filler { ptr };
21
22
          cgh.parallel_for(r, filler);
23
        });
24 }
```

4.14.2 Defining kernels as lambda functions

In C++, function objects can be defined using lambda functions. Kernels may be defined as lambda functions in SYCL. The name of a lambda function in SYCL may optionally be specified by passing it as a template parameter to the invoking member function, and in that case, the lambda name is a C++ typename. If the lambda function relies on template arguments, then if specified, the name of the lambda function must contain those template arguments. The class used for the name of a lambda function is only used for naming purposes and is not required to be defined. For details on restrictions for kernel naming, please refer to 5.2.

The kernel function for the lambda function is the lambda function itself. The kernel lambda must use copy for all of its captures (i.e. [=]), and the lambda must not use the mutable specifier.

```
1 class MyKernel;
2
3 myQueue.submit([&](handler& cmdGroup) {
4 cmdgroup.single_task<class MyKernel>([=]() {
5 // [kernel code]
6 });
7 });
```

4.14.3 Defining kernels using modules

In case the developer needs to specify compiler flags or special linkage options for a kernel, then a kernel object can be used, as described in 4.13.6.2. The SYCL kernel function is defined as a named function object 4.14.1 or lambda function 4.14.2. The user can obtain a module object for the kernel with the get_kernel member function. This member function is templated by the kernel name, so that the user can specify the kernel whose associated kernel they wish to obtain.

The following example illustrates how an application can pre-compile a kernel. The code defines the kernel as a lambda function, and pre-compiles it by creating a module object for all the kernels defined in the current translation unit. This ensures that the kernel is compiled *a priori* of its invocation via parallel_for. For more details on the module object and its related APIs, see Section 4.13.

```
1
   sycl::queue myQueue;
 2
    auto myContext = myQueue.get_context();
3
4
   // Calling the get() function will pre-compile all kernels in this translation
5
   // unit for the device in "myContext".
6
    auto myModule = sycl::this_module::get<sycl::module_state::executable>(myContext);
7
8
    auto myRange = sycl::nd_range<2>(range<2>(1024, 1024),range<2>(64, 64));
9
10
   myQueue.submit([&](sycl::handler& cgh) {
     // Calling use_module() causes the parallel_for() below to use the pre-compiled
11
      // kernel from "myModule".
12
13
      cgh.use_module(myModule);
14
15
      cgh.parallel_for(myRange, ([=](sycl::nd_item<2> index) {
        // kernel code
16
17
     }));
18 });
```

In the following example, the SYCL kernel function performs a convolution and uses specialization constants to set the values of the coefficients.

```
#include <CL/sycl.hpp>
1
2
3
   using namespace sycl;
4
5
   using coeff_t = std::array<std::array<float, 3>, 3>;
6
7
    // Read coefficients from somewhere.
8
   coeff_t get_coefficients();
9
10
   // Identify the specialization constant.
    specialization_id<coeff_t> coeff_id;
11
12
13
    void do_conv(buffer<float, 2> in, buffer<float, 2> out) {
14
     queue myQueue;
15
16
      myQueue.submit([&](handler &cgh) {
17
        accessor in_acc { in, cgh, read_only };
18
        accessor out_acc { out, cgh, write_only };
```

```
19
        // Set the coefficient of the convolution as constant.
20
21
        // This will build a specific kernel the coefficient available as literals.
22
        cgh.set_specialization_constant<coeff_id>(get_coefficients());
23
24
        cgh.parallel_for<class Convolution>(
25
            in.get_range(), [=](item<2> item_id, kernel_handler h) {
26
              float acc = 0;
27
              coeff_t coeff = h.get_specialization_constant<coeff_id>();
28
              for (int i = -1; i \le 1; i ++) {
29
                if (item_id[0] + i < 0 || item_id[0] + i >= in_acc.get_range()[0])
30
                  continue:
31
                for (int j = -1; j <= 1; j++) {
32
                  if (item_id[1] + j < 0 || item_id[1] + j >= in_acc.get_range()[1])
33
                    continue:
34
                  // the underlying JIT can see all the values of the array returned by coeff.get().
35
                  acc += coeff[i + 1][j + 1] *
36
                         in_acc[item_id[0] + i][item_id[1] + j];
37
                }
38
              }
39
              out_acc[item_id] = acc;
40
            });
41
      });
42
43
      myQueue.wait();
44 }
```

4.14.4 Rules for parameter passing to kernels

In a case where a kernel is a named function object or a lambda function, any member variables encapsulated within the function object or variables captured by the lambda function must be treated according to the following rules:

- Any accessor must be passed as an argument to the device kernel in a form that allows the device kernel to access the data in the specified way.
- The SYCL runtime and compiler(s) must produce the necessary conversions to enable accessor arguments from the host to be converted to the correct type of parameter on the device.
- The device compiler(s) must validate that the layout of any data shared between the host and the device(s) (e.g. value kernel arguments or data accessed through an accessor or USM) is compatible with the layout of that data on the host. If there is a layout mismatch that the implementation cannot or will not correct for (to make the layouts compatible), then the device compiler must issue an error and compilation must fail.
- A local accessor provides access to work-group-local memory. The accessor is not constructed with any buffer, but instead constructed with a size and base data type. The runtime must ensure that the work-group-local memory is allocated per work-group and available to be used by the kernel via the local accessor.
- C++ trivially copyable types must be passed by value to the kernel.
- C++ non-trivially copyable types must not be passed as arguments to a kernel that is compiled for a device.
- It is illegal to pass a buffer or image (instead of an accessor class) as an argument to a kernel. Generation of a compiler error in this illegal case is optional.

- Sampler objects (sampler) can be passed as parameters to kernels.
- It is illegal to pass a pointer or reference argument to a kernel. Generation of a compiler error in this illegal case is optional.
- Any aggregate types such as structs or classes should follow the rules above recursively. It is not necessary to separate struct or class members into separate kernel parameters if all members of the aggregate type are unaffected by the rules above.

4.15 Error handling

4.15.1 Error handling rules

Error handling in a SYCL application (host code) uses C++ exceptions. If an error occurs, it will be thrown by the API function call and may be caught by the user through standard C++ exception handling mechanisms.

SYCL applications are asynchronous in the sense that host and device code executions are decoupled from one another except at specific points. For example, device code executions often begin when dependencies in the SYCL task graph are satisfied, which occurs asynchronously from host code execution. As a result of this the errors that occur on a device cannot be thrown directly from a host API call, because the call enqueueing a device action has typically already returned by the time that the error occurs. Such errors are not detected until the error-causing task executes or tries to execute, and we refer to these as asynchronous errors.

4.15.1.1 Asynchronous error handler

The queue and context classes can optionally take an asynchronous handler object <code>async_handler</code> on construction, which is a callable such as a function class or lambda, with an <code>exception_list</code> as a parameter. Invocation of an <code>async_handler</code> may be triggered by the queue member functions <code>queue::wait_and_throw()</code> or <code>queue:: throw_asynchronous()</code>, by the event member function <code>event::wait_and_throw()</code>, or automatically on destruction of a queue or context that contains unconsumed asynchronous errors. When invoked, an <code>async_handler</code> is called and receives an <code>exception_list</code> argument containing a list of exception objects representing any unconsumed asynchronous errors associated with the queue or context.

When an asynchronous error instance has been passed to an async_handler, then that instance of the error has been consumed for handling and is not reported on any subsequent invocations of the async_handler.

The async_handler may be a named function object type, a lambda function or a std::function. The exception_list object passed to the async_handler is constructed by the SYCL runtime.

4.15.1.2 Behavior without an async_handler

If an asynchronous error occurs in a queue or context that has no user-supplied asynchronous error handler object async_handler, then an implementation defined default async_handler is called to handle the error in the same situations that a user-supplied async_handler would be, as defined in 4.15.1.1. The default async_handler must in some way report all errors passed to it, when possible, and must then invoke std::terminate or equivalent.

4.15.1.3 Priorities of async_handlers

If the SYCL runtime can associate an asynchronous error with a specific queue, then:

• If the queue was constructed with an async_handler, that handler is invoked to handle the error.

- Otherwise if the context enclosed by the queue was constructed with an async_handler, that handler is invoked to handle the error.
- Otherwise when no handler was passed to either queue or context on construction, then a default handler is invoked to handle the error, as described by 4.15.1.2.
- All handler invocations in this list occur at times as defined by 4.15.1.1.

If the SYCL runtime cannot associate an asynchronous error with a specific queue, then:

- If the context in which the error occurred was constructed with an async_handler, then that handler is invoked to handle the error.
- Otherwise when no handler was passed to the associated context on construction, then a default handler is invoked to handle the error, as described by 4.15.1.2.
- All handler invocations in this paragraph occur at times as defined by 4.15.1.1.

4.15.1.4 Asynchronous errors with a secondary queue

If an asynchronous error occurs when running or enqueuing a command group which has a secondary queue specified, then the command group may be enqueued to the secondary queue instead of the primary queue. The error handling in this case is also configured using the async_handler provided for both queues. If there is no async_handler given on any of the queues, then no asynchronous error reporting is done and no exceptions are thrown. If the primary queue fails and there is an async_handler given at this queue's construction, which populates the exception_list parameter, then any errors will be added and can be thrown whenever the user chooses to handle those exceptions. Since there were errors on the primary queue and a secondary queue was given, then the execution of the kernel is re-scheduled to the secondary queue and any error reporting for the kernel execution on that queue is done through that queue, in the same way as described above. The secondary queue may fail as well, and the errors will be thrown if there is an async_handler and either wait_and_throw() or throw() are called on that queue. The command group function object event returned by that function will be relevant to the queue where the kernel has been enqueued.

Below is an example of catching a SYCL exception and printing out the error message.

```
1 void catch_any_errors(sycl::context const& ctx) {
2   try {
3     do_something_to_invoke_error(ctx);
4   }
5     catch(sycl::exception const& e) {
6     std::cerr << e.what();
7   }
8 }</pre>
```

Below is an example of catching a SYCL exception with the invalid_object error code and printing out the error message.

```
1 void catch_invalid_object_errors(sycl::context const& ctx) {
2  try {
3   do_something_to_invoke_error(ctx);
4   }
5   catch(sycl::exception const& e) {
```

Below is an example of catching a SYCL exception, checking for the SYCL backend by inspecting the category and handling the OpenCL SYCL backend error codes if the category is that of the OpenCL SYCL backend otherwise checking the standard error code.

```
void catch_backend_errors(sycl::context const& ctx) {
 1
2
      try {
3
        do_something_to_invoke_error(ctx);
4
      }
5
      catch(sycl::exception const& e) {
6
        if(e.category() == sycl::error_category_for<sycl::backend::opencl>()) {
7
          switch(e.code().value()) {
8
            case CL_INVALID_PROGRAM:
9
              std::cerr << "OpenCL invalid program error: " << e.what();</pre>
10
            /* ...*/
          }
11
12
          else {
13
            throw;
14
          }
15
        }
16
        else {
17
          if(e.code() == sycl::errc::invalid_object) {
            std::cerr << "Invalid object error: " << e.what();</pre>
18
19
          }
20
          else {
21
            throw;
22
          }
23
        }
24
      }
25 }
```

4.15.2 Exception class interface

```
1 namespace sycl {
2
3 using async_handler = std::function<void(sycl::exception_list)>;
4
5 class exception : public virtual std::exception {
6
    public:
7
       exception(std::error_code ec, const std::string& what_arg);
8
       exception(std::error_code ec, const char * what_arg);
9
       exception(std::error_code ec);
10
       exception(int ev, const std::error_category& ecat, const std::string& what_arg);
11
       exception(int ev, const std::error_category& ecat, const char* what_arg);
```

```
12
        exception(int ev, const std::error_category& ecat);
13
14
        exception(context ctx, std::error_code ec, const std::string& what_arg);
15
        exception(context ctx, std::error_code ec, const char* what_arg);
16
        exception(context ctx, std::error_code ec);
17
        exception(context ctx, int ev, const std::error_category& ecat, const std::string& what_arg);
18
        exception(context ctx, int ev, const std::error_category& ecat, const char* what_arg);
19
        exception(context ctx, int ev, const std::error_category& ecat);
20
21
        const std::error_code& code() const noexcept;
22
        const std::error_category& category() const noexcept;
23
24
        bool has_context() const noexcept;
25
        context get_context() const;
26 };
27
28 class exception_list {
29
     // Used as a container for a list of asynchronous exceptions
30
     public:
31
     using value_type = std::exception_ptr;
32
      using reference = value_type&;
33
      using const_reference = const value_type&;
34
      using size_type = std::size_t;
35
      using iterator = /*unspecified*/;
      using const_iterator = /*unspecified*/;
36
37
38
      size_type size() const;
39
      iterator begin() const; // first asynchronous exception
40
      iterator end() const; // refer to past-the-end last asynchronous exception
41 };
42
43 enum class errc {
      runtime_error = /* implementation-defined */,
44
45
      kernel = /* implementation-defined */,
      accessor = /* implementation-defined */,
46
      nd_range = /* implementation-defined */,
47
      event = /* implementation-defined */,
48
49
      invalid_parameter = /* implementation-defined */,
      compile_program = /* implementation-defined */,
50
51
      link_program = /* implementation-defined */,
52
      invalid_object = /* implementation-defined */,
53
      memory_allocation = /* implementation-defined */,
54
      platform = /* implementation-defined */,
55
      profiling = /* implementation-defined */,
      feature_not_supported = /* implementation-defined */
56
57 };
58
59 template<backend b>
    using errc_for = typename backend_traits<b>::errc;
60
61
62 std::error_condition make_error_condition(errc e) noexcept;
63
    std::error_code make_error_code(errc e) noexcept;
64
65
    const std::error_category& sycl_category() noexcept;
66
```
```
template<backend b>
67
    const std::error_category& error_category_for() noexcept;
68
69
70
   } // namespace sycl
71
72
   namespace std {
73
74
     template <>
75
     struct is_error_condition_enum<sycl::errc> : true_type {};
76
77
      template <>
78
      struct is_error_code_enum<see-below> : true_type {};
79
80 } // namespace std
```

The SYCL exception_list class is also available in order to provide a list of synchronous and asynchronous exceptions.

Errors can occur both in the SYCL library and SYCL host side, or may come directly from a SYCL backend. The member functions on these exceptions provide the corresponding information. SYCL backends can provide additional exception class objects as long as they derive from sycl::exception object, or any of its derived classes.

A specialization of std::is_error_condition_enum must be defined for sycl::errc inheriting from std:: true_type.

A specialization of std::is_error_code_enum must be defined for sycl::errc and backend_traits<Backend>:: errc inheriting from std::true_type for each Backend, where backend is each enumeration of the enum class backend.

| Member function | Description |
|---|--|
| <pre>exception(std::error_code ec, const std::string&</pre> | Constructs an exception. The string re- |
| what_arg) | turned by what() is guaranteed to contain |
| | what_arg as a substring. |
| <pre>exception(std::error_code ec, const char* what_arg)</pre> | Constructs an exception. The string re- |
| | turned by what() is guaranteed to contain |
| | what_arg as a substring. |
| <pre>exception(std::error_code ec)</pre> | Constructs an exception. |
| <pre>exception(int ev, const std::error_category& ecat,</pre> | Constructs an exception with the error code |
| <pre>const std::string& what_arg)</pre> | ev and the underlying error category ecat. |
| | The string returned by what () is guaranteed |
| | to contain what_arg as a substring. |
| <pre>exception(int ev, const std::error_category& ecat,</pre> | Constructs an exception with the error code |
| <pre>const char* what_arg)</pre> | ev and the underlying error category ecat. |
| | The string returned by what () is guaranteed |
| | to contain what_arg as a substring. |
| <pre>exception(int ev, const std::error_category& ecat)</pre> | Constructs an exception with the error code |
| | ev and the underlying error category ecat. |
| | Continued on next page |

Table 4.101: Member functions of the SYCL exception class.

| Member function | Description |
|---|--|
| <pre>exception(context ctx, std::error_code ec, const std</pre> | Constructs an exception with an associated |
| ::string& what_arg) | SYCL context ctx. The string returned by |
| | what() is guaranteed to contain what_arg as |
| | a substring. |
| <pre>exception(context ctx, std::error_code ec, const</pre> | Constructs an exception with an associated |
| <pre>char* what_arg)</pre> | SYCL context ctx. The string returned by |
| | what() is guaranteed to contain what_arg as |
| | a substring. |
| <pre>exception(context ctx, std::error_code ec)</pre> | Constructs an exception with an associated |
| | SYCL context ctx. |
| <pre>exception(context ctx, int ev, const std::</pre> | Constructs an exception with an associated |
| error_category& ecat, const std::string& what_arg) | SYCL context ctx, the error code ev and the |
| | underlying error category ecat. The string |
| | returned by what () is guaranteed to contain |
| | what_arg as a substring. |
| <pre>exception(context ctx, int ev, const std::</pre> | Constructs an exception with an associated |
| error_category& ecat, const char* what_arg) | SYCL context ctx, the error code ev and the |
| | underlying error category ecat. The string |
| | returned by what () is guaranteed to contain |
| | what_arg as a substring. |
| <pre>exception(context ctx, int ev, const std::</pre> | Constructs an exception with an associated |
| error_category& ecat) | SYCL context ctx, the error code ev and the |
| | underlying error category ecat. |
| <pre>const std::error_code& code()const noexcept</pre> | Returns the error code stored inside the ex- |
| | ception. |
| <pre>const std::error_categeory& category()const noexcept</pre> | Returns the error category of the error code |
| | stored inside the exception. |
| <pre>const char *what()const</pre> | Returns an implementation defined non-null |
| | constant C-style string that describes the er- |
| | ror that triggered the exception. |
| <pre>bool has_context()const</pre> | Returns true if this SYCL exception has |
| | an associated SYCL context and false if |
| | it does not. |
| <pre>context get_context()const</pre> | Returns the SYCL context that is associated |
| | with this SYCL exception if one is avail- |
| | able. Must throw an exception with the |
| | errc::invalid_object_error error code if |
| | this SYCL exception does not have a SYCL |
| | context. |
| | End of table |

Table 4.101: Member functions of the SYCL exception class.

| Member function | Description |
|----------------------------------|---|
| <pre>size_t size()const</pre> | Returns the size of the list |
| <pre>iterator begin()const</pre> | Returns an iterator to the beginning of the |
| | list of asynchronous exceptions. |
| | Continued on next page |

Table 4.102: Member functions of the exception_list.

| Member function | Description |
|---------------------|---|
| iterator end()const | Returns an iterator to the end of the list of |
| | asynchronous exceptions. |
| | End of table |

Table 4.102: Member functions of the exception_list.

| Standard SYCL Error Codes | Description |
|------------------------------------|---|
| runtime_error | Generic runtime error. |
| kernel_error | Error that occurred before or while enqueu- |
| | ing the SYCL kernel. |
| nd_range_error | Error regarding the SYCL nd_range speci- |
| | fied for the SYCL kernel |
| accessor_error | Error regarding the SYCL accessor objects |
| | defined. |
| event_error | Error regarding associated SYCL event ob- |
| | jects. |
| invalid_parameter_error | Error regarding parameters to the SYCL ker- |
| | nel, it may apply to any captured parameters |
| | to the kernel lambda. |
| compile_program_error | Error while compiling the SYCL kernel to a |
| | SYCL device. |
| link_program_error | Error while linking the SYCL kernel to a |
| | SYCL device. |
| invalid_object_error | Error regarding any memory objects being |
| | used inside the kernel |
| <pre>memory_allocation_error</pre> | Error on memory allocation on the SYCL |
| | device for a SYCL kernel. |
| platform_error | The SYCL platform will trigger this excep- |
| | tion on error. |
| profiling_error | The SYCL runtime will trigger this error if |
| | there is an error when profiling info is en- |
| | abled. |
| <pre>feature_not_supported</pre> | Exception thrown when an optional feature |
| | or extension is used in a kernel but it's not |
| | available on the device the SYCL kernel is |
| | being enqueued on. |
| | End of table |

Table 4.103: Values of the SYCL errc enum.

| SYCL Error Code Helpers | Description |
|--|---|
| <pre>std::error_condition make_error_condition(errc e)</pre> | Constructs an error condition using e and |
| noexcept; | <pre>sycl_category().</pre> |
| <pre>std::error_code make_error_code(errc e)noexcept;</pre> | Constructs an error code using e and |
| | <pre>sycl_category().</pre> |
| | End of table |

Table 4.104: SYCL error code helper functions.

4.16 Data types

SYCL as a C++ programming model supports the C++ core language data types, and it also provides the ability for all SYCL applications to be executed on SYCL compatible devices. The scalar and vector data types that are supported by the SYCL system are defined below. More details about the SYCL device compiler support for fundamental and OpenCL interoperability types are found in 5.5.

4.16.1 Scalar data types

The fundamental C++ data types which are supported in SYCL are described in Table 5.1. Note these types are fundamental and therefore do not exist within the sycl namespace.

Additional scalar data types which are supported by SYCL within the sycl namespace are described in Table 4.105.

| Scalar data type | Description |
|------------------|--|
| byte | An unsigned 8-bit integer. This is depre- |
| | cated in SYCL 2020 since C++17 std:: |
| | byte can be used instead. |
| half | A 16-bit floating-point. The half data |
| | type must conform to the IEEE 754-2008 |
| | half precision storage format. A SYCL |
| | <pre>feature_not_supported exception must be</pre> |
| | thrown if the half type is used in a SYCL |
| | kernel function which executes on a SYCL |
| | device that does not have aspect::fp16. |
| | End of table |

Table 4.105: Additional scalar data types supported by SYCL.

4.16.2 Vector types

SYCL provides a cross-platform class template that works efficiently on SYCL devices as well as in host C++ code. This type allows sharing of vectors between the host and its SYCL devices. The vector supports member functions that allow construction of a new vector from a swizzled set of component elements.

vec<typename dataT, int numElements> is a vector type that compiles down to a SYCL backend built-in vector types on SYCL devices, where possible, and provides compatible support on the host or when it is not possible. The vec class is templated on its number of elements and its element type. The number of elements parameter, numElements, can be one of: 1, 2, 3, 4, 8 or 16. Any other value shall produce a compilation failure. The element type parameter, dataT, must be one of the basic scalar types supported in device code.

The SYCL vec class template provides interoperability with the underlying vector type defined by vector_t which is available only when compiled for the device. The SYCL vec class can be constructed from an instance of vector_t and can implicitly convert to an instance of vector_t in order to support interoperability with native SYCL backend functions from a SYCL kernel function.

An instance of the SYCL vec class template can also be implicitly converted to an instance of the data type when the number of elements is 1 in order to allow single element vectors and scalars to be convertible with each other.

4.16.2.1 Vec interface

The constructors, member functions and non-member functions of the SYCL vec class template are listed in Tables 4.106, 4.107 and 4.108 respectively.

```
namespace sycl {
 1
 2
 3
    enum class rounding_mode {
 4
      automatic,
 5
     rte,
 6
     rtz,
 7
     rtp,
 8
      rtn
 9
    };
10
11 struct elem {
12
     static constexpr int x = 0;
13
    static constexpr int y = 1;
14
    static constexpr int z = 2;
15
    static constexpr int w = 3;
16
    static constexpr int r = 0;
17
    static constexpr int g = 1;
18
     static constexpr int b = 2;
19
     static constexpr int a = 3;
20
     static constexpr int s0 = 0;
21
     static constexpr int s1 = 1;
22
     static constexpr int s2 = 2;
23
     static constexpr int s3 = 3;
24
     static constexpr int s4 = 4;
25
     static constexpr int s5 = 5;
26
     static constexpr int s6 = 6;
27
     static constexpr int s7 = 7;
28
     static constexpr int s8 = 8;
29
     static constexpr int s9 = 9;
30
     static constexpr int sA = 10;
31
     static constexpr int sB = 11;
32
     static constexpr int sC = 12;
33
      static constexpr int sD = 13;
34
      static constexpr int sE = 14;
35
      static constexpr int sF = 15;
36 };
37
    template <typename dataT, int numElements>
38
39
    class vec {
40
    public:
41
     using element_type = dataT;
42
43
    #ifdef __SYCL_DEVICE_ONLY__
44
      using vector_t = __unspecified__;
45
    #endif
46
47
      vec();
48
49
      explicit vec(const dataT &arg);
50
```

```
51
       template <typename... argTN>
52
       vec(const argTN&... args);
53
54
      vec(const vec<dataT, numElements> &rhs);
55
56
    #ifdef __SYCL_DEVICE_ONLY__
57
      vec(vector_t openclVector);
58
59
      operator vector_t() const;
60
    #endif
61
      // Available only when: numElements == 1
62
63
      operator dataT() const;
64
65
      static constexpr int get_count();
66
67
      static constexpr size_t get_size();
68
69
      template <typename convertT, rounding_mode roundingMode = rounding_mode::automatic>
70
       vec<convertT, numElements> convert() const;
71
72
      template <typename asT>
73
      asT as() const;
74
75
       template<int... swizzleIndexes>
76
      __swizzled_vec__ swizzle() const;
77
78
      // Available only when numElements <= 4.</pre>
79
      // XYZW_ACCESS is: x, y, z, w, subject to numElements.
80
      __swizzled_vec__ XYZW_ACCESS() const;
81
82
      // Available only numElements == 4.
83
      // RGBA_ACCESS is: r, g, b, a.
84
      __swizzled_vec__ RGBA_ACCESS() const;
85
86
      // INDEX_ACCESS is: s0, s1, s2, s3, s4, s5, s6, s7, s8, s9, sA, sB, sC, sD,
87
      // sE, sF, subject to numElements.
88
      __swizzled_vec__ INDEX_ACCESS() const;
89
90 #ifdef SYCL_SIMPLE_SWIZZLES
91
      // Available only when numElements <= 4.</pre>
92
      // XYZW_SWIZZLE is all permutations with repetition of: x, y, z, w, subject to
93
      // numElements.
94
      95
96
      // Available only when numElements == 4.
97
      // RGBA_SWIZZLE is all permutations with repetition of: r, g, b, a.
98
      __swizzled_vec__ RGBA_SWIZZLE() const;
99
100
    #endif // #ifdef SYCL_SIMPLE_SWIZZLES
101
      // Available only when: numElements > 1.
102
103
      __swizzled_vec__ lo() const;
104
      __swizzled_vec__ hi() const;
105
      __swizzled_vec__ odd() const;
```

```
106
       __swizzled_vec__ even() const;
107
108
       // load and store member functions
109
       template <access::address_space addressSpace, access::decorated IsDecorated>
110
       void load(size_t offset, multi_ptr<const dataT, addressSpace, IsDecorated> ptr);
111
       template <access::address_space addressSpace, access::decorated IsDecorated>
112
       void store(size_t offset, multi_ptr<dataT, addressSpace, IsDecorated> ptr) const;
113
114
       // subscript operator
115
       dataT &operator[](int index);
       const dataT &operator[](int index) const;
116
117
118
       // OP is: +, -, *, /, %
119
       /* If OP is % available, only when: dataT != float && dataT != double
120
       && dataT != half. */
121
       friend vec operatorOP(const vec &lhs, const vec &rhs) { /* ... */ }
122
       friend vec operatorOP(const vec &lhs, const dataT &rhs) { /* ... */ }
123
124
       // OP is: +=, -=, *=, /=, %=
125
       /* If OP is %= available, only when: dataT != float && dataT != double
126
       && dataT != half. */
       friend vec &operatorOP(vec &lhs, const vec &rhs) { /* ... */ }
127
128
       friend vec &operatorOP(vec &lhs, const dataT &rhs) { /* ... */ }
129
130
       // OP is: ++, --
131
       friend vec &operatorOP(vec &lhs) { /* ... */ }
132
       friend vec operatorOP(vec& lhs, int) { /* ... */ }
133
134
       // OP is: +, -
135
       friend vec operatorOP(vec &lhs) const { /* ... */ }
136
137
      // OP is: &, |, ^
138
       /* Available only when: dataT != float && dataT != double
139
       && dataT != half. */
       friend vec operatorOP(const vec &lhs, const vec &rhs) { /* ... */ }
140
141
       friend vec operatorOP(const vec &lhs, const dataT &rhs) { /* ... */ }
142
       // OP is: &=, |=, ^=
143
144
       /* Available only when: dataT != float && dataT != double
145
       && dataT != half. */
       friend vec &operatorOP(vec &lhs, const vec &rhs) { /* ... */ }
146
147
       friend vec &operatorOP(vec &lhs, const dataT &rhs) { /* ... */ }
148
149
       // OP is: &&, ||
150
       friend vec<RET, numElements> operatorOP(const vec &lhs, const vec &rhs) { /* ... */ }
151
       friend vec<RET, numElements> operatorOP(const vec& lhs, const dataT &rhs) { /* ... */ }
152
153
      // OP is: <<, >>
154
       /* Available only when: dataT != float && dataT != double
155
       && dataT != half. */
156
       friend vec operatorOP(const vec &lhs, const vec &rhs) { /* ... */ }
157
       friend vec operatorOP(const vec &lhs, const dataT &rhs) { /* ... */ }
158
159
       // OP is: <<=, >>=
160
       /* Available only when: dataT != float && dataT != double
```

```
161
       && dataT != half. */
       friend vec &operatorOP(vec &lhs, const vec &rhs) { /* ... */ }
162
163
       friend vec &operatorOP(vec &lhs, const dataT &rhs) { /* ... */ }
164
165
       // OP is: ==, !=, <, >, <=, >=
166
       friend vec<RET, numElements> operatorOP(const vec &lhs, const vec &rhs) { /* ... */ }
167
       friend vec<RET, numElements> operatorOP(const vec &lhs, const dataT &rhs) { /* ... */ }
168
169
       vec &operator=(const vec<dataT, numElements> &rhs);
       vec &operator=(const dataT &rhs);
170
171
172
       /* Available only when: dataT != float && dataT != double
173
       && dataT != half. */
174
       friend vec operator~(const vec &v) { /* ... */ }
175
       friend vec<RET, numElements> operator!(const vec &v) { /* ... */ }
176
177
       // OP is: +, -, *, /, %
178
       /* operator% is only available when: dataT != float && dataT != double &&
179
       dataT != half. */
180
       friend vec operatorOP(const dataT &lhs, const vec &rhs) { /* ... */ }
181
       // OP is: &, |, ^
182
       /* Available only when: dataT != float && dataT != double
183
       && dataT != half. */
184
185
       friend vec operatorOP(const dataT &lhs, const vec &rhs) { /* ... */ }
186
187
       // OP is: &&, ||
       friend vec<RET, numElements> operatorOP(const dataT &lhs, const vec &rhs) { /* ... */ }
188
189
190
       // OP is: <<, >>
191
       /* Available only when: dataT != float && dataT != double
192
      && dataT != half. */
193
       friend vec operatorOP(const dataT &lhs, const vec &rhs) { /* ... */ }
194
195
      // OP is: ==, !=, <, >, <=, >=
196
       friend vec<RET, numElements> operatorOP(const dataT &lhs, const vec &rhs) { /* ... */ }
197
198 };
199
200 // Deduction guides
201 // Available only when: (std::is_same_v<T, U> && ...)
202 template <class T, class... U>
203 vec(T, U...) -> vec<T, sizeof...(U) + 1>;
204
205 } // namespace sycl
```

| Constructor | Description |
|-------------|---|
| vec() | Default construct a vector with element type |
| | dataT and with numElements dimensions by |
| | default construction of each of its elements. |
| | Continued on next page |

Table 4.106: Constructors of the SYCL vec class template.

| Constructor | Description |
|--|---|
| <pre>explicit vec(const dataT &arg)</pre> | Construct a vector of element type dataT |
| | and numElements dimensions by setting |
| | each value to arg by assignment. |
| <pre>template <typename argtn=""></typename></pre> | Construct a SYCL vec instance from any |
| <pre>vec(const argTN& args)</pre> | combination of scalar and SYCL vec param- |
| | eters of the same element type, providing the |
| | total number of elements for all parameters |
| | sum to numElements of this vec specializa- |
| | tion. |
| <pre>vec(const vec<datat, numelements=""> &rhs)</datat,></pre> | Construct a vector of element type dataT |
| | and number of elements numElements by |
| | copy from another similar vector. |
| <pre>vec(vector_t openclVector)</pre> | Available only when: compiled for the de- |
| | vice. |
| | Constructs a SYCL vec instance from an |
| | instance of the underlying OpenCL vector |
| | type defined by vector_t. |
| | End of table |

Table 4.106: Constructors of the SYCL vec class template.

| Member function | Description |
|---|---|
| operator vector_t()const | Available only when: compiled for the de- |
| | vice. |
| | Converts this SYCL vec instance to the un- |
| | derlying OpenCL vector type defined by |
| | vector_t. |
| operator dataT()const | Available only when: numElements == 1. |
| | Converts this SYCL vec instance to an in- |
| | stance of dataT with the value of the single |
| | element in this SYCL vec instance. |
| | The SYCL vec instance shall be implicitly |
| | convertible to the same data types, to which |
| | dataT is implicitly convertible. Note that |
| | conversion operator shall not be templated to |
| | allow standard conversion sequence for im- |
| | plicit conversion. |
| <pre>static constexpr int get_count()</pre> | Returns the number of elements of this |
| | SYCL vec. |
| <pre>static constexpr size_t get_size()</pre> | Returns the size of this SYCL vec in bytes. |
| | 3-element vector size matches 4-element |
| | vector size to provide interoperability with |
| | OpenCL vector types. The same rule applies |
| | to vector alignment as described in 4.16.2.6. |
| | Continued on next page |

Table 4.107: Member functions for the SYCL vec class template.

| Member function | Description |
|--|---|
| <pre>template<typename convertt,="" pre="" rounding_mode<=""></typename></pre> | Converts this SYCL vec to a SYCL vec |
| <pre>roundingMode = rounding_mode::automatic></pre> | of a different element type specified by |
| <pre>vec<convertt, numelements=""> convert()const</convertt,></pre> | convertT using the rounding mode speci- |
| | fied by roundingMode. The new SYCL vec |
| | type must have the same number of elements |
| | as this SYCL vec. The different rounding |
| | modes are described in Table 4.109. |
| <pre>template<typename ast=""></typename></pre> | Bitwise reinterprets this SYCL vec as a |
| asT as()const | SYCL vec of a different element type and |
| | number of elements specified by asT. The |
| | new SYCL vec type must have the same |
| | storage size in bytes as this SYCL vec. |
| <pre>template<int swizzleindexes=""></int></pre> | Return an instance of the implementa- |
| <pre>swizzled_vec swizzle()const</pre> | tion defined intermediate class template |
| | swizzled_vec representing an index se- |
| | quence which can be used to apply the swiz- |
| | zle in a valid expression as described in |
| | 4.16.2.4. |
| swizzled_vec XYZW_ACCESS()const | Available only when: numElements <= 4. |
| | Returns an instance of the implementa- |
| | tion defined intermediate class template |
| | swizzled_vec representing an index |
| | sequence which can be used to apply the |
| | swizzle in a valid expression as described in |
| | 4.10.2.4. |
| | Where XYZW ACCESS is: x for numElements |
| | -1 y y for numElements -2 y y 7 |
| | for numElements -3 and \mathbf{x} \mathbf{y} \mathbf{z} w for |
| | numFlements $= 4$ |
| swizzled vec RGRA ACCESS()const | Available only when: numFlements == 4 |
| | Returns an instance of the implementa- |
| | tion defined intermediate class template |
| | swizzled vec representing an index |
| | sequence which can be used to apply the |
| | swizzle in a valid expression as described in |
| | 4.16.2.4. |
| | |
| | Where RGBA_ACCESS is: r, g, b, a. |
| | Continued on next page |

| Member function | Description |
|----------------------------------|--|
| swizzled_vec INDEX_ACCESS()const | Returns an instance of the implementa- tion defined intermediate class template swizzled_vec representing an index sequence which can be used to apply the swizzle in a valid expression as described in 4.16.2.4. |
| | Where INDEX_ACCESS is: s0 for numElements == 1, s0, s1 for numElements == 2, s0, s1, s2 for numElements == 3, s0 , s1, s2, s3 for numElements == 4, s0, s1, s2, s3, s4, s5, s6, s7, s8 for numElements == 8 and s0, s1, s2, s3, s4, s5, s6, s7, s8, s9, sA, sB, sC, sD, sE, sF for numElements == 16. |
| swizzled_vec XYZW_SWIZZLE()const | Available only when numElements <= 4, and when the macro SYCL_SIMPLE_SWIZZLES is defined before including sycl.hpp. Returns an instance of the implementa- tion defined intermediate class template swizzled_vec representing an index sequence which can be used to apply the swizzle in a valid expression as described in 4.16.2.4. |
| | Where XYZW_SWIZZLE is all per- mutations with repetition, of any subset with length greater than 1, of x, y for numElements == 2, x, y, z for numElements == 3 and x, y, z, w for numElements == 4. For example a four el- ement vec provides permutations including xzyw, xyyy and xz. |
| swizzled_vec RGBA_SWIZZLE()const | Available only when numElements == 4, and when the macro SYCL_SIMPLE_SWIZZLES is defined before including sycl.hpp. Returns an instance of the implementa- tion defined intermediate class template swizzled_vec representing an index sequence which can be used to apply the swizzle in a valid expression as described in 4.16.2.4. |
| | Where RGBA_SWIZZLE is all permu- tations with repetition, of any subset with length greater than 1, of r, g, b, a. For example a four element vec provides permutations including rbga, rggg and rb. Continued on next page |

| Member function | Description |
|--|---|
| swizzled_vec lo()const | Available only when: numElements > 1. |
| | Return an instance of the implementa- |
| | tion defined intermediate class template |
| | swizzled_vec representing an index se- |
| | quence made up of the lower half of this |
| | SYCL vec which can be used to apply the |
| | swizzle in a valid expression as described |
| | in $4.16.2.4$. When numElements == 3, this |
| | SYCL vec is treated as though numElements |
| | == 4 with the fourth element undefined. |
| swizzled_vec hi()const | Available only when: numElements > 1 . |
| | Return an instance of the implementa- |
| | tion defined intermediate class template |
| | swizzled_vec representing an index se- |
| | quence made up of the upper half of this |
| | SYCL vec which can be used to apply the |
| | swizzle in a valid expression as described |
| | in 4.16.2.4. When numElements $=$ 3, this |
| | SYCL vec is treated as though numElements |
| | == 4 with the fourth element undefined. |
| swizzled_vec odd()const | Available only when: numElements > 1 . |
| | Return an instance of the implementa- |
| | tion defined intermediate class template |
| | swizzled_vec representing an index se- |
| | quence made up of the odd indexes of this |
| | SYCL vec which can be used to apply the |
| | in 4.16.2.4. When sumElements 2. this |
| | 1114.10.2.4. When numerements == 5, this SVCL upped is tracted as though numElements |
| | A with the fourth element undefined |
| guizzled wee even() const | |
| | Available only when numerements > 1 . Return an instance of the implementa |
| | tion defined intermediate class template |
| | swizzled vec representing an index se- |
| | guence made up of the even indexes of this |
| | SYCL vec which can be used to apply the |
| | swizzle in a valid expression as described |
| | in 4.16.2.4. When numElements $=$ 3, this |
| | SYCL vec is treated as though numElements |
| | == 4 with the fourth element undefined. |
| <pre>template <access::address_space access<="" addressspace,="" pre=""></access::address_space></pre> | Loads the values at the address of ptr offset |
| ::decorated IsDecorated> | in elements of type dataT by numElements * |
| <pre>void load(size_t offset, multi_ptr<const datat,<="" pre=""></const></pre> | offset, into the components of this SYCL |
| addressSpace, IsDecorated> ptr) | vec. |
| <pre>template <access::address_space access<="" addressspace,="" pre=""></access::address_space></pre> | Stores the components of this SYCL vec |
| ::decorated IsDecorated> | into the values at the address of ptr offset |
| <pre>void store(size_t offset, multi_ptr<datat,< pre=""></datat,<></pre> | in elements of type dataT by numElements |
| addressSpace, IsDecorated> ptr)const | * offset. |
| | Continued on next page |

| Member function | Description |
|--|---|
| <pre>dataT &operator[](int index)</pre> | Returns a reference to the element stored |
| | within this SYCL vec at the index specified |
| | by index. |
| <pre>const dataT &operator[](int index)const</pre> | Returns a const reference to the element |
| | stored within this SYCL vec at the index |
| | specified by index. |
| <pre>vec &operator=(const vec &rhs)</pre> | Assign each element of the rhs SYCL vec |
| | to each element of this SYCL vec and return |
| | a reference to this SYCL vec. |
| <pre>vec &operator=(const dataT &rhs)</pre> | Assign each element of the rhs scalar to |
| | each element of this SYCL vec and return |
| | a reference to this SYCL vec. |
| | End of table |

| vec operatorOP(const vec &lhs, const vec &rhs) If OP is % available, only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL vec class template with the same template parameters as lhs vec with each element of the new SYCL vec instance the result of an element-wise OP arithmetic operation between each element of the rhs SYCL vec. Where OP is: +, -, *, /, %. vec operatorOP(const vec &lhs, const dataT &rhs) If OP is % available, only when: dataT != float && dataT != half. Construct a new instance of the SYCL vec. Where OP is: +, -, *, /, %. vec operatorOP(const vec &lhs, const dataT &rhs) If OP is % available, only when: dataT != float && dataT != half. Construct a new instance of the SYCL vec class template with the same template parameters as lhs vec with each element of the new SYCL vec instance the result of an element-wise OP arithmetic operation between each element of lhs vec and the rhs scalar. Where OP is: +, -, *, /, %. | Hidden friend function | Description |
|---|---|---|
| Where OP is: +, -, *, /, %.vec operatorOP(const vec &lhs, const dataT &rhs)If OP is % available, only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL vec class template with the same template parameters as lhs vec with each element of the new SYCL vec instance the result of an element-wise OP arithmetic operation between each element of lhs vec and the rhs scalar. Where OP is: +, -, *, /, %. | <pre>vec operatorOP(const vec &lhs, const vec &rhs)</pre> | If OP is % available, only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL vec class template with the same template parameters as 1hs vec with each element of the new SYCL vec instance the result of an element-wise OP arithmetic operation between each element of 1hs vec and each element of the rhs SYCL vec. |
| vec operatorOP(const vec &lhs, const dataT &rhs)If OP is % available, only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL vec class template with the same template parameters as 1hs vec with each element of the new SYCL vec instance the result of an element-wise OP arithmetic operation between each element of 1hs vec and the rhs scalar. Where OP is: +, -, *, /, %. | | Where OP is: +, -, *, /, %. |
| Where OP 1s: +, -, *, /, %. | <pre>vec operatorOP(const vec &lhs, const dataT &rhs)</pre> | If OP is % available, only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL vec class template with the same template parameters as lhs vec with each element of the new SYCL vec instance the result of an element-wise OP arithmetic operation between each element of lhs vec and the rhs scalar. |
| Continued on next needs | | Where OP 1s: +, -, *, /, %. |

| Hidden friend function | Description |
|--|---|
| <pre>vec &operatorOP(vec &lhs, const vec &rhs)</pre> | If OP is %= available, only when: dataT != float && dataT != double && dataT != half. Perform an in-place element-wise OP arith- metic operation between each element of lhs vec and each element of the rhs SYCL vec and return lhs vec. |
| <pre>vec &operatorOP(vec &lhs, const dataT &rhs)</pre> | Where OP is: +=, -=, *=, /=, %=. If OP is %= available, only when: dataT != float && dataT != double && dataT != half. Perform an in-place element-wise OP arithmetic operation between each element of lhs vec and rhs scalar and return lhs vec. Where OP is: += -= *= /= %= |
| vec &operatorOP(vec &v) | Perform an in-place element-wise OP prefix arithmetic operation on each element of 1hs vec, assigning the result of each element to the corresponding element of 1hs vec and return 1hs vec. |
| <pre>vec operatorOP(vec &v, int)</pre> | Where OP is: ++, Perform an in-place element-wise OP post- fix arithmetic operation on each element of 1hs vec, assigning the result of each element to the corresponding element of 1hs vec and returns a copy of 1hs vec before the operation is performed. Where OP is: ++, |
| <pre>vec operatorOP(vec &v)</pre> | Construct a new instance of the SYCL vec class template with the same template parameters as this SYCL vec with each element of the new SYCL vec instance the result of an element-wise OP unary arithmetic operation on each element of this SYCL vec. |
| | Where OP is: +, Continued on next page |
| | Continued on next page |

| Hidden friend function | Description |
|---|--|
| <pre>vec operatorOP(const vec &lhs, const vec &rhs)</pre> | Available only when: dataT != float && |
| | dataT != double && dataT != half. |
| | Construct a new instance of the SYCL |
| | vec class template with the same template |
| | parameters as 1hs vec with each element of |
| | the new SYCL vec instance the result of an |
| | element-wise OP bitwise operation between |
| | each element of 1hs vec and each element |
| | of the rhs SYCL vec. |
| | |
| | Where OP is: &, , ^. |
| <pre>vec operatorOP(const vec &lhs, const dataT &rhs)</pre> | Available only when: dataT != float && |
| | <pre>dataT != double && dataT != half.</pre> |
| | Construct a new instance of the SYCL |
| | vec class template with the same template |
| | parameters as 1hs vec with each element of |
| | the new SYCL vec instance the result of an |
| | element-wise OP bitwise operation between |
| | each element of 1hs vec and the rhs scalar. |
| | |
| | Where OP is: &, , ^. |
| <pre>vec &operatorOP(vec &lhs, const vec &rhs)</pre> | Available only when: dataT != float && |
| | dataT != double && dataT != half. |
| | Perform an in-place element-wise OP bitwise |
| | operation between each element of 1hs vec |
| | and the rhs SYCL vec and return 1hs vec. |
| | |
| | Where OP is: &=, =, ^=. |
| <pre>vec &operatorOP(vec &lhs, const dataT &rhs)</pre> | Available only when: dataT != float && |
| | dataT != double && dataT != half. |
| | Perform an in-place element-wise OP bitwise |
| | operation between each element of 1hs vec |
| | and the rhs scalar and return a lhs vec. |
| | William op in a la c |
| | Where OP 1s: &=, =, ^=. |
| | Continued on next page |

| <pre>vec<ret, numelements=""> operatorOP(const vec &lhs, const vec &rhs)</ret,></pre> Construct a new instance of the SYCL vec class template with the same template parameters as hs vec with each element of the new SYCL vec instance the result of an element-wise OP logical operation between each element of the rws SYCL vec. and each element of the rhs SYCL vec. The dataT template parameter of this SYCL vec. For a SYCL vec with dataT of type int8_t or uint8_t RET must be int8_t. For a SYCL vec with dataT of type int16_t, uint16_t or half RET must be int64_t. Where OP is: &&, . Vec <ret, numelements=""> operatorOP(const vec &lhs, const dataT &rhs) Vec<ket, numelements=""> operatorOP(const vec &lhs, const dataT &rhs) Construct a new instance of the SYCL vec instance of the synce intact the same template parameters as this SYCL vec with dataT of type int64_t, uint64_t or double RET must be int64_t. The dataT template parameter of the constructed SYCL vec, with each element of the new SYCL vec instance the result of an element.vise OP logical operation between each element of the synce and the rhs scalar. The dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of the constructed SYCL vec, with dataT of type int6_t, tuin16_t or half RET must be int8_t. For a SYCL vec with dataT of type int6_t, tuin16_t or float RET must be int6_t.t. Where OP is: &k, . Continued on part means the the start of the the the the the the the the the the</ket,></ret,> | Hidden friend function | Description |
|--|---|--|
| The dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of this SYCL vec. For a SYCL vec with dataT of type int8_t or uint8_t RET must be int8_t. For a SYCL vec with dataT of type int16_t, uint16_t or half RET must be int32_t, uint32_t or float RET must be int64_t.vec <ret, numelements=""> operatorOP(const vec &lhs, const dataT & rhs)Construct a new instance of the SYCL vec with dataT of type int64_t, uint64_t or double RET must be int64_t.vec<ret, numelements=""> operatorOP(const vec &lhs, const dataT & rhs)Construct a new instance of the SYCL vec instance the result of an element-wise OP logical operation between each element of lbs vec and the rhs scalar.The dataT template parameter of this SYCL vec. For a SYCL vec with dataT of type int64_t, uint64_t or double RET must be int32_t. For a SYCL vec with dataT of type int64_t.the dataT template parameter of the constructed SYCL vec. RET, varies depend- ing on the dataT template parameter of this SYCL vec. For a SYCL vec with dataT of type int64_t. uint64_t or double RET must be int16_t. For a SYCL vec with dataT of type int16_t. Ther a SYCL vec with dataT of type int64_t. uint64_t or double RET must be int16_t. For a SYCL vec with dataT of type int64_t. uint64_t or double RET must be int16_t. For a SYCL vec with dataT of type int64_t. uint64_t or double RET must be int64_t.</ret,></ret,> | <pre>vec<ret, numelements=""> operatorOP(const vec &lhs, const vec &rhs)</ret,></pre> | Construct a new instance of the SYCL vec class template with the same template parameters as 1hs vec with each element of the new SYCL vec instance the result of an element-wise OP logical operation between each element of 1hs vec and each element of the rhs SYCL vec. |
| Vec <ret, numelements=""> operatorOP(const vec &lhs, const dataT &rhs)Construct a new instance of the SYCL vec class template with the same template parameters as this SYCL vec with each element of the new SYCL vec instance the result of an element-wise OP logical operation between each element of lhs vec and the rhs scalar.The dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of this SYCL vec. For a SYCL vec with dataT of type int16_t, uint16_t or half RET must be int16_t, uint16_t or float RET must be int16_t, uint64_t or double RET must be int164_t.Where OP is: &&, .</ret,> | | The dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of this SYCL vec. For a SYCL vec with dataT of type int8_t or uint8_t RET must be int8_t. For a SYCL vec with dataT of type int16_t, uint16_t or half RET must be int16_t. For a SYCL vec with dataT of type int32_t, uint32_t or float RET must be int32_t. For a SYCL vec with dataT of type int64_t, uint64_t or double RET must be int64_t. |
| The dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of this SYCL vec. For a SYCL vec with dataT of type int8_t or uint8_t RET must be int8_t. For a SYCL vec with dataT of type int16_t, uint16_t or half RET must be int16_t. For a SYCL vec with dataT of type int32_t, uint32_t or float RET must be int32_t. For a SYCL vec with dataT of type int64_t, uint64_t or double RET must be uint64_t. Where OP is: &&, . | <pre>vec<ret, numelements=""> operatorOP(const vec &lhs, const dataT &rhs)</ret,></pre> | Where OP is: &&, . Construct a new instance of the SYCL vec class template with the same template parameters as this SYCL vec with each element of the new SYCL vec instance the result of an element-wise OP logical operation between each element of 1hs vec and the rhs scalar. |
| be int32_t. For a SYCL vec with dataT of type int64_t, uint64_t or double RET must be uint64_t. Where OP is: &&, . | | The dataT template parameter of the constructed SYCL vec, RET, varies depending on the dataT template parameter of this SYCL vec. For a SYCL vec with dataT of type int8_t or uint8_t RET must be int8_t. For a SYCL vec with dataT of type int16_t, uint16_t or half RET must be int16_t. For a SYCL vec with dataT of type int32_t, uint32_t or float RET must |
| Where OP is: &&, . | | be int32_t. For a SYCL vec with dataT of type int64_t, uint64_t or double RET must be uint64_t. |
| | | Where OP is: &&, . |

| Hidden friend function | Description |
|---|---|
| Hidden friend function vec operatorOP(const vec &lhs, const vec &rhs) | Description Available only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL vec class template with the same template parameters as lhs vec with each element of the new SYCL vec instance the result of an element-wise OP bitshift operation between each element of lhs vec and each element of the rhs SYCL vec. If OP is >>, dataT is a signed type and lhs vec has a negative value any vacated bits viewed as an unsigned integer must be assigned the value 1, otherwise any vacated bits viewed as an unsigned integer must be assigned the value 0. Where OP is: << >> |
| <pre>vec operatorOP(const vec &lhs, const dataT &rhs)</pre> | Available only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL vec class template with the same template parameters as 1hs vec with each element of the new SYCL vec instance the result of an element-wise OP bitshift operation between each element of 1hs vec and the rhs scalar. If OP is >>, dataT is a signed type and 1hs vec has a negative value any vacated bits viewed as an unsigned integer must be assigned the value 1, otherwise any vacated bits viewed as an unsigned integer must be assigned the value 0. |
| <pre>vec &operatorOP(vec &lhs, const vec &rhs)</pre> | Where OP is: <<, >>. Available only when: dataT != float && dataT != double && dataT != half. Perform an in-place element-wise OP bitshift operation between each element of 1hs vec and the rhs SYCL vec and returns 1hs vec. If OP is >>=, dataT is a signed type and 1hs vec has a negative value any vacated bits viewed as an unsigned integer must be assigned the value 1, otherwise any vacated bits viewed as an unsigned integer must be assigned the value 0. |
| | Continued on next page |

| Hidden friend function | Description |
|---|---|
| <pre>vec &operatorOP(vec &lhs, const dataT &rhs)</pre> | Available only when: dataT != float && dataT != double && dataT != half. Perform an in-place element-wise OP bit- shift operation between each element of lhs vec and the rhs scalar and returns a reference to this SYCL vec. If OP is >>=, dataT is a signed type and lhs vec has a negative value any vacated bits viewed as an unsigned integer must be assigned the value 1, otherwise any vacated bits viewed as an unsigned integer must be assigned the value 0. |
| | Where OP is: <<=, >>=. |
| <pre>vec<ret, numelements=""> operatorOP(const vec& lhs, const vec &rhs)</ret,></pre> | Construct a new instance of the SYCL vec class template with the element type RET with each element of the new SYCL vec instance the result of an element-wise OP relational operation between each element of 1hs vec and each element of the rhs SYCL vec. Each element of the SYCL vec that is returned must be -1 if the operation results in true and 0 if the operation results in false or either this SYCL vec or the rhs SYCL vec is a NaN. |
| | The dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of this SYCL vec. For a SYCL vec with dataT of type int8_t or uint8_t RET must be int8_t. For a SYCL vec with dataT of type int16_t, uint16_t or half RET must be int16_t. For a SYCL vec with dataT of type int32_t, uint32_t or float RET must be int32_t. For a SYCL vec with dataT of type int64_t, uint64_t or double RET must be uint64_t. |
| | Where OP is: ==, !=, <, >, <=, >=. |

| Hidden friend function | Description |
|---|---|
| <pre>vec<ret, numelements=""> operatorOP(const vec &lhs,</ret,></pre> | Construct a new instance of the SYCL vec |
| const dataT &rhs) | class template with the dataT parameter of |
| | RET with each element of the new SYCL vec |
| | instance the result of an element-wise OP |
| | relational operation between each element |
| | of lhs vec and the rhs scalar. Each element |
| | of the SYCL vec that is returned must be -1 |
| | if the operation results in true and 0 if the |
| | operation results in false or either lhs vec |
| | or the rhs SYCL vec is a NaN. |
| | The dataT template parameter of the |
| | constructed SVCL vec PET varies depend- |
| | ing on the dataT template parameter of this |
| | SYCL vec For a SYCL vec with data |
| | of type int8 t or uint8 t RFT must be |
| | int8 t. For a SYCL vec with dataT of |
| | type int16 t. uint16 t or half RET must |
| | be int16_t. For a SYCL vec with dataT of |
| | type int32_t, uint32_t or float RET must |
| | be int32_t. For a SYCL vec with dataT of |
| | type int64_t, uint64_t or double RET must |
| | be uint64_t. |
| | |
| | Where OP is: ==, !=, <, >, <=, >=. |
| <pre>vec operatorOP(const dataT &lhs, const vec &rhs)</pre> | If OP is % available, only when: dataT != |
| | <pre>float && dataT != double && dataT !=</pre> |
| | half. |
| | Construct a new instance of the SYCL |
| | vec class template with the same template |
| | element of the new SVCL was instance |
| | the result of an element-wise OP arithmetic |
| | operation between the lbs scalar and each |
| | element of the rhs SYCL vec. |
| | |
| | Where OP is: +, -, *, /, %. |
| <pre>vec operatorOP(const dataT &lhs, const vec &rhs)</pre> | Available only when: dataT != float && |
| | <pre>dataT != double && dataT != half.</pre> |
| | Construct a new instance of the SYCL |
| | vec class template with the same template |
| | parameters as the rhs SYCL vec with each |
| | the result of an element wise OD bitwise |
| | operation between the like cooler and each |
| | element of the rbs SVCL was |
| | Clement of the fills 5 f CL vec. |
| | Where OP is: &, , ^. |
| | Continued on next page |

| Hidden friend function | Description |
|---|--|
| <pre>vec<ret, numelements=""> operatorOP(const dataT &lhs,</ret,></pre> | Available only when: dataT != float && |
| const vec &rhs) | <pre>dataT != double && dataT != half.</pre> |
| | Construct a new instance of the SYCL |
| | vec class template with the same template |
| | parameters as the rhs SYCL vec with each |
| | element of the new SYCL vec instance |
| | the result of an element-wise OP logical |
| | operation between the 1hs scalar and each |
| | element of the rhs SYCL vec. |
| | |
| | The dataT template parameter of the |
| | constructed SYCL vec, RET, varies depend- |
| | ing on the dataT template parameter of this |
| | SYCL vec. For a SYCL vec with dataT |
| | of type int8_t or uint8_t RET must be |
| | int8_t. For a SYCL vec with dataT of |
| | type intl6_t, uint16_t of half REI must |
| | be inti6_t. For a SYCL vec with datal of |
| | be int22 + For a SVCL use with dataT of |
| | two interferences interferences and the set of double PET must |
| | be int64_t, ullit04_t of double RET must |
| | |
| | Where OP is: &&, . |
| <pre>vec operatorOP(const dataT &lhs, const vec &rhs)</pre> | Construct a new instance of the SYCL |
| | vec class template with the same template |
| | parameters as the rhs SYCL vec with each |
| | element of the new SYCL vec instance |
| | the result of an element-wise OP bitshift |
| | operation between the 1hs scalar and each |
| | element of the rhs SYCL vec. If OP is >>, |
| | dataT is a signed type and this SYCL vec |
| | nas a negative value any vacated bits viewed |
| | as an unsigned integer must be assigned the |
| | value 1, otherwise any vacaled bits viewed |
| | as an unsigned integer must be assigned the |
| | |
| | Where OP is: <<. >>. |
| | Continued on next page |

| Hidden friend function | Description |
|---|--|
| <pre>vec<ret, numelements=""> operatorOP(const dataT &lhs, const vec &rhs)</ret,></pre> | Available only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL vec class template with the element type RET with each element of the new SYCL vec instance the result of an element-wise OP relational operation between the 1hs scalar and each element of the rhs SYCL vec. Each element of the SYCL vec that is returned must be -1 if the operation results in true and 0 if the operation results in false or either this SYCL vec or the rhs SYCL vec is a NaN. |
| | The dataT template parameter of the constructed SYCL vec, RET, varies depending on the dataT template parameter of this SYCL vec. For a SYCL vec with dataT of type int8_t or uint8_t RET must be int8_t. For a SYCL vec with dataT of type int16_t, uint16_t or half RET must be int16_t. For a SYCL vec with dataT of type int32_t, uint32_t or float RET must be int32_t. For a SYCL vec with dataT of type int64_t, uint64_t or double RET must be int64_t. |
| | Where OP is: ==, !=, <, >, <=, >=. |
| vec &operator~(const vec &v) | Available only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL vec class template with the same template pa- rameters as v vec with each element of the new SYCL vec instance the result of an element-wise OP bitwise operation on each element of v vec. |
| | Continued on next page |

| Hidden friend function | Description |
|--|--|
| <pre>vec<ret, numelements=""> operator!(const vec &v)</ret,></pre> | Construct a new instance of the SYCL vec class template with the same template parameters as v vec with each element of the new SYCL vec instance the result of an element-wise OP logical operation on each element of v vec. Each element of the SYCL vec that is returned must be -1 if the operation results in true and 0 if the operation results in false or this SYCL vec is a NaN. |
| | The dataT template parameter of the constructed SYCL vec, RET, varies depend- ing on the dataT template parameter of this SYCL vec. For a SYCL vec with dataT of type int8_t or uint8_t RET must be int8_t. For a SYCL vec with dataT of type int16_t, uint16_t or half RET must be int16_t. For a SYCL vec with dataT of type int32_t, uint32_t or float RET must be int32_t. For a SYCL vec with dataT of type int64_t, uint64_t or double RET must be int64_t. |
| | End of table |

4.16.2.2 Aliases

The SYCL programming API provides all permutations of the type alias:

using <type><elems> = vec<<storage-type>, <elems>>

where <elems> is 2, 3, 4, 8 and 16, and pairings of <type> and <storage-type> for integral types are char and int8_t, uchar and uint8_t, short and int16_t, ushort and uint16_t, int and int32_t, uint and uint32_t, long and int64_t, ulong and uint64_t, and for floating point types are both half, float and double.

For example uint4 is the alias to vec<uint64_t, 4> and float16 is the alias to vec<float, 16>.

4.16.2.3 Swizzles

Swizzle operations can be performed in two ways. Firstly by calling the swizzle member function template, which takes a variadic number of integer template arguments between 0 and numElements-1, specifying swizzle indexes. Secondly by calling one of the simple swizzle member functions defined in 4.107 as XYZW_SWIZZLE and RGBA_SWIZZLE. Note that the simple swizzle functions are only available for up to 4 element vectors and are only available when the macro SYCL_SIMPLE_SWIZZLES is defined before including sycl.hpp.

In both cases the return type is always an instance of __swizzled_vec__, an implementation defined temporary class representing a swizzle of the original SYCL vec instance. Both kinds of swizzle member functions must not perform the swizzle operation themselves, instead the swizzle operation must be performed by the returned

instance of __swizzled_vec__ when used within an expression, meaning if the returned __swizzled_vec__ is never used in an expression no swizzle operation is performed.

Both the swizzle member function template and the simple swizzle member functions allow swizzle indexes to be repeated.

A series of static constexpr values are provided within the elem struct to allow specifying named swizzle indexes when calling the swizzle member function template.

4.16.2.4 Swizzled vec class

The __swizzled_vec__ class must define an unspecified temporary which provides the entire interface of the SYCL vec class template, including swizzled member functions, with the additions and alterations described below:

- The __swizzled_vec__ class template must be readable as an r-value reference on the RHS of an expression. In this case the swizzle operation is performed on the RHS of the expression and then the result is applied to the LHS of the expression.
- The __swizzled_vec__ class template must be assignable as an l-value reference on the LHS of an expression. In this case the RHS of the expression is applied to the original SYCL vec which the __swizzled_vec__ represents via the swizzle operation. Note that a __swizzled_vec__ that is used in an l-value expression may not contain any repeated element indexes. For example: f4.xxxx()= fx.wzyx() would not be valid.
- The __swizzled_vec__ class template must be convertible to an instance of SYCL vec with the type dataT and number of elements specified by the swizzle member function, if numElements > 1, and must be convertible to an instance of type dataT, if numElements == 1.
- The __swizzled_vec__ class template must be non-copyable, non-moveable, non-user constructible and may not be bound to a l-value or escape the expression it was constructed in. For example auto x = f4.x () would not be valid.
- The __swizzled_vec__ class template should return __swizzled_vec__ & for each operator inhetired from the vec class template interface which would return vec<dataT, numElements> &.

4.16.2.5 Rounding modes

The various rounding modes that can be used in the as member function template are described in Table 4.109.

| Rounding mode | Description |
|---------------|---|
| automatic | Default rounding mode for the SYCL vec |
| | class element type. rtz (round toward zero) |
| | for integer types and rte (round to nearest |
| | even) for floating-point types. |
| rte | Round to nearest even. |
| rtz | Round toward zero. |
| rtp | Round toward positive infinity. |
| rtn | Round toward negative infinity. |
| | End of table |

Table 4.109: Rounding modes for the SYCL vec class template.

4.16.2.6 Memory layout and alignment

The elements of an instance of the SYCL vec class template are stored in memory sequentially and contiguously and are aligned to the size of the element type in bytes multiplied by the number of elements:

The exception to this is when the number of element is three in which case the SYCL vec is aligned to the size of the element type in bytes multiplied by four:

$$sizeof(dataT) \cdot 4$$
 (4.7)

This is true for both host and device code in order to allow for instances of the vec class template to be passed to SYCL kernel functions.

4.16.2.7 Considerations for endianness

As SYCL supports both big-endian and little-endian on OpenCL devices, users must take care to ensure kernel arguments are processed correctly. This is particularly true for SYCL vec arguments as the order in which a SYCL vec is loaded differs between big-endian and little-endian.

Users should consult vendor documentation for guidance on how to handle kernel arguments in these situations.

4.16.2.8 Performance note

The usage of the subscript operator[] may not be efficient on some devices.

4.16.3 Math array types

SYCL provides a marray<typename dataT, std::size_t numElements> class template to represent a contiguous fixed-size container. This type allows sharing of containers between the host and its SYCL devices.

The marray class is templated on its element type and number of elements. The number of elements parameter, numElements, is a positive value of the std::size_t type. The element type parameter, dataT, must be a Numeric type as it is defined by C++ standard.

An instance of the marray class template can also be implicitly converted to an instance of the data type when the number of elements is 1 in order to allow single element arrays and scalars to be convertible with each other.

Logical and comparison operators for marray class template return marray
bool, numElements>.

4.16.3.1 Math array interface

The constructors, member functions and non-member functions of the SYCL marray class template are listed in Tables 4.110, 4.111 and 4.112 respectively.

```
1 namespace sycl {
2
3 template <typename dataT, std::size_t numElements>
4 class marray {
```

```
5
     public:
 6
      using value_type = dataT;
 7
      using reference = dataT&;
 8
      using const_reference = const dataT&;
 9
      using iterator = dataT*;
10
      using const_iterator = const dataT*;
11
12
     marray();
13
14
      explicit marray(const dataT &arg);
15
16
      template <typename... argTN>
17
      marray(const argTN&... args);
18
19
      marray(const marray<dataT, numElements> &rhs);
20
      marray(marray<dataT, numElements> &&rhs);
21
22
      // Available only when: numElements == 1
23
      operator dataT() const;
24
25
      static constexpr std::size_t size();
26
27
      // subscript operator
28
      reference operator[](std::size_t index);
29
      const_reference operator[](std::size_t index) const;
30
      marray &operator=(const marray<dataT, numElements> &rhs);
31
32
     marray &operator=(const dataT &rhs);
33 };
34
35 // iterator functions
36 iterator begin(marray &v);
37
   const_iterator begin(const marray &v);
38
39 iterator end(marray &v);
40 const_iterator end(const marray &v);
41
42 // OP is: +, -, *, /, %
   /* If OP is % available, only when: dataT != float && dataT != double && dataT != half. */
43
44 marray operatorOP(const marray &lhs, const marray &rhs) { /* ... */ }
45 marray operatorOP(const marray &lhs, const dataT &rhs) { /* ... */ }
46
47 // OP is: +=, -=, *=, /=, %=
48 /* If OP is %= available, only when: dataT != float && dataT != double && dataT != half. */
49 marray &operatorOP(marray &lhs, const marray &rhs) { /* ... */ }
50 marray &operatorOP(marray &lhs, const dataT &rhs) { /* ... */ }
51
52 // OP is: ++, --
53 marray &operatorOP(marray &lhs) { /* ... */ }
54 marray operatorOP(marray& lhs, int) { /* ... */ }
55
56 // OP is: +, -
57 marray operatorOP(marray &lhs) const { /* ... */ }
58
59 // OP is: &, |, ^
```

```
60 /* Available only when: dataT != float && dataT != double && dataT != half. */
61 marray operatorOP(const marray &lhs, const marray &rhs) { /* ... */ }
62 marray operatorOP(const marray &lhs, const dataT &rhs) { /* ... */ }
63
64 // OP is: &=, |=, ^=
65 /* Available only when: dataT != float && dataT != double && dataT != half. */
66 marray &operatorOP(marray &lhs, const marray &rhs) { /* ... */ }
67 marray & operatorOP(marray & lhs, const dataT & rhs) { /* ... */ }
68
69 // OP is: &&, ||
70 marray<bool, numElements> operatorOP(const marray &lhs, const marray &rhs) { /* ... */ }
71 marray<bool, numElements> operatorOP(const marray& lhs, const dataT &rhs) { /* ... */ }
72
73 // OP is: <<, >>
74 /* Available only when: dataT != float && dataT != double && dataT != half. */
75 marray operatorOP(const marray &lhs, const marray &rhs) { /* ... */ }
76 marray operatorOP(const marray &lhs, const dataT &rhs) { /* ... */ }
77
78 // OP is: <<=, >>=
79 /* Available only when: dataT != float && dataT != double && dataT != half. */
80 marray &operatorOP(marray &lhs, const marray &rhs) { /* ... */ }
81 marray &operatorOP(marray &lhs, const dataT &rhs) { /* ... */ }
82
83 // OP is: ==, !=, <, >, <=, >=
84 marray<br/>
whool, numElements> operatorOP(const marray &lhs, const marray &rhs) { /* ... */ }
85 marray<bool, numElements> operatorOP(const marray &lhs, const dataT &rhs) { /* ... */ }
86
87 /* Available only when: dataT != float && dataT != double && dataT != half. */
88 marray operator~(const marray &v) { /* ... */ }
89 marray<bool, numElements> operator!(const marray &v) { /* ... */ }
90
91 // OP is: +, -, *, /, %
92 /* operator% is only available when: dataT != float && dataT != double && dataT != half. */
93 marray operatorOP(const dataT &lhs, const marray &rhs) { /* ... */ }
94
95 // OP is: &, |, ^
96 /* Available only when: dataT != float && dataT != double
97 && dataT != half. */
98 marray operatorOP(const dataT &lhs, const marray &rhs) { /* ... */ }
99
100 // OP is: &&, ||
101 marray<bool, numElements> operatorOP(const dataT &lhs, const marray &rhs) { /* ... */ }
102
103 // OP is: <<, >>
104 /* Available only when: dataT != float && dataT != double && dataT != half. */
105 marray operatorOP(const dataT &lhs, const marray &rhs) { /* ... */ }
106
107 // OP is: ==, !=, <, >, <=, >=
108 marray<bool, numElements> operatorOP(const dataT &lhs, const marray &rhs) { /* ... */ }
109
110 marray operator~(const marray &v) const { /* ... */ }
111
112 marray<bool, numElements> operator!(const marray &v) const { /* ... */ }
113
114 } // namespace sycl
```

| Constructor | Description |
|--|--|
| marray() | Default construct an array with element |
| | type dataT and with numElements dimen- |
| | sions by default construction of each of its |
| | elements. |
| <pre>explicit marray(const dataT &arg)</pre> | Construct an array of element type dataT |
| | and numElements dimensions by setting |
| | each value to arg by assignment. |
| <pre>template <typename argtn=""></typename></pre> | Construct a SYCL marray instance from any |
| <pre>marray(const argTN& args)</pre> | combination of scalar and SYCL marray pa- |
| | rameters of the same element type, provid- |
| | ing the total number of elements for all pa- |
| | rameters sum to numElements of this marray |
| | specialization. |
| <pre>marray(const marray<datat, numelements=""> &rhs)</datat,></pre> | Construct an array of element type dataT |
| | and number of elements numElements by |
| | copy from another similar vector. |
| | End of table |

Table 4.110: Constructors of the SYCL marray class template.

| Member function | Description |
|--|---|
| operator dataT()const | Available only when: numElements == 1. |
| | Converts this SYCL marray instance to an |
| | instance of dataT with the value of the sin- |
| | gle element in this SYCL marray instance. |
| | The SYCL marray instance shall be im- |
| | plicitly convertible to the same data types, |
| | to which dataT is implicitly convertible. |
| | Note that conversion operator shall not be |
| | templated to allow standard conversion se- |
| | quence for implicit conversion. |
| <pre>static constexpr std::size_t size()</pre> | Returns the size of this SYCL marray in |
| | bytes. |
| | 3-element vector size matches 4-element |
| | vector size to provide interoperability with |
| | OpenCL vector types. The same rule applies |
| | to vector alignment as described in 4.16.2.6. |
| <pre>dataT &operator[](std::size_t index)</pre> | Returns a reference to the element stored |
| | within this SYCL marray at the index speci- |
| | fied by index. |
| <pre>const dataT &operator[](std::size_t index)const</pre> | Returns a const reference to the element |
| | stored within this SYCL marray at the index |
| | specified by index. |
| <pre>marray &operator=(const marray &rhs)</pre> | Assign each element of the rhs SYCL |
| | marray to each element of this SYCL marray |
| | and return a reference to this SYCL marray. |
| | Continued on next page |

Table 4.111: Member functions for the SYCL marray class template.

| Member function | Description |
|--|--|
| <pre>marray &operator=(const dataT &rhs)</pre> | Assign each element of the rhs scalar to |
| | each element of this SYCL marray and re- |
| | turn a reference to this SYCL marray. |
| | End of table |

Table 4.111: Member functions for the SYCL marray class template.

| Non-member function | Description |
|--|---|
| iterator begin(marray &v) | Returns an iterator referring to the first ele- |
| | ment stored within the v marray. |
| <pre>const_iterator begin(const marray &v)</pre> | Returns a const iterator referring to the first |
| | element stored within the v marray. |
| iterator end(marray &v) | Returns an iterator referring to the one past |
| | the last element stored within the v marray. |
| <pre>const_iterator end(const marray &v)</pre> | Returns a const iterator referring to the one |
| | past the last element stored within the v |
| | marray. |
| <pre>marray operatorOP(const marray &lhs, const marray &</pre> | If OP is % available, only when: dataT != |
| rhs) | <pre>float && dataT != double && dataT !=</pre> |
| | half. |
| | Construct a new instance of the SYCL |
| | marray class template with the same tem- |
| | plate parameters as 1hs marray with each |
| | element of the new SYCL marray instance |
| | the result of an element-wise OP arithmetic |
| | operation between each element of 1hs |
| | marray and each element of the rhs SYCL |
| | marray. |
| | |
| | Where OP is: +, -, *, /, %. |
| marray operatorOP(const marray &Ihs, const dataT & | If OP is % available, only when: data1 != |
| rhs) | float && datal != double && datal != |
| | half. |
| | Construct a new instance of the SYCL |
| | marray class template with the same tem- |
| | element of the new SVCI manney instance |
| | the result of an element wise OP orithmetic |
| | operation between each element of the |
| | marray and the rhs scalar |
| | marray and the rus scalar. |
| | Where OP is: $+ - * / \%$ |
| | Continued on next page |

Table 4.112: Non-member functions of the marray class template.

| Non-member function | Description |
|---|---|
| <pre>marray &operatorOP(marray &lhs, const marray &rhs)</pre> | If OP is %= available, only when: dataT != |
| | <pre>float && dataT != double && dataT !=</pre> |
| | half. |
| | Perform an in-place element-wise OP arith- |
| | metic operation between each element of |
| | lhs marray and each element of the rhs |
| | SYCL marray and return lhs marray. |
| | |
| | Where OP 1s: +=, -=, *=, /=, %=. |
| <pre>marray &operatorOP(marray &lhs, const dataT &rhs)</pre> | If OP is %= available, only when: dataT != |
| | <pre>float && dataT != double && dataT !=</pre> |
| | halt. |
| | Perform an in-place element-wise op arith- |
| | the mean and the scalar and return the |
| | The marray and the scalar and feturin the |
| | Mallay. |
| | Where OP is: $+= -= *= /= \%=$ |
| marray & operator (P(marray & y) | Perform an in-place element-wise OP prefix |
| marray asperators (marray av) | arithmetic operation on each element of 1hs |
| | marray, assigning the result of each element |
| | to the corresponding element of lhs marray |
| | and return lhs marray. |
| | |
| | Where OP is: ++, |
| <pre>marray operatorOP(marray &v, int)</pre> | Perform an in-place element-wise OP post- |
| | fix arithmetic operation on each element |
| | of lhs marray, assigning the result of each |
| | element to the corresponding element of 1hs |
| | marray and returns a copy of lhs marray |
| | before the operation is performed. |
| | Where OP is: |
| marray operator (P(marray & y) | Construct a new instance of the SVCI |
| marray operatorol (marray av) | marray class template with the same tem- |
| | plate parameters as this SYCL marray with |
| | each element of the new SYCL marray |
| | instance the result of an element-wise OP |
| | unary arithmetic operation on each element |
| | of this SYCL marray. |
| | |
| | Where OP is: +, |
| | Continued on next page |

| Non-member function | Description |
|---|---|
| <pre>marray operatorOP(const marray &lhs, const marray & rhs)</pre> | Available only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL marray class template with the same tem- plate parameters as 1hs marray with each element of the new SYCL marray instance the result of an element-wise OP bitwise operation between each element of 1hs marray and each element of the rhs SYCL marray. |
| <pre>marray operatorOP(const marray &lhs, const dataT & rhs)</pre> | Available only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL marray class template with the same tem- plate parameters as 1hs marray with each element of the new SYCL marray instance the result of an element-wise OP bitwise operation between each element of 1hs marray and the rhs scalar. |
| <pre>marray &operatorOP(marray &lhs, const marray &rhs)</pre> | Where OP is: &, , ^. Available only when: dataT != float && dataT != double && dataT != half. Perform an in-place element-wise OP bit- wise operation between each element of lhs marray and the rhs SYCL marray and return lhs marray. Where OP is: &= l= ^= |
| marray &operatorOP(marray &lhs, const dataT &rhs) | Where OP is: &=, =, ^=. Available only when: dataT != float && dataT != double && dataT != half. Perform an in-place element-wise OP bit- wise operation between each element of lhs marray and the rhs scalar and return a lhs marray. Where OP is: &=, l=, ^=. |
| <pre>marray<bool, numelements=""> operatorOP(const marray & lhs, const marray &rhs)</bool,></pre> | Construct a new instance of the marray class template with dataT = bool and same numElements as 1hs marray with each element of the new marray instance the result of an element-wise OP logical operation between each element of 1hs marray and each element of the rhs marray. Where OP is: &&, . |
| | Continued on next page |

| Non-member function | Description |
|--|---|
| <pre>marray<bool, numelements=""> operatorOP(const marray & lhs, const dataT &rhs)</bool,></pre> | Construct a new instance of the marray class template with dataT = bool and same numElements as lhs marray with each element of the new marray instance the result of an element-wise OP logical operation between each element of lhs marray and the rhs scalar. |
| | Where OP is: &&, 11. |
| <pre>marray operatorOP(const marray &lhs, const marray & rhs)</pre> | Available only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL marray class template with the same tem- plate parameters as 1hs marray with each element of the new SYCL marray instance the result of an element-wise OP bitshift operation between each element of 1hs marray and each element of the rhs SYCL marray. If OP is >>, dataT is a signed type and 1hs marray has a negative value any vacated bits viewed as an unsigned integer must be assigned the value 1, otherwise any vacated bits viewed as an unsigned integer must be assigned the value 0. |
| | Where OP is: <<, >>. |
| <pre>marray operatorOP(const marray &lhs, const dataT & rhs)</pre> | Available only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL marray class template with the same tem- plate parameters as 1hs marray with each element of the new SYCL marray instance the result of an element-wise OP bitshift operation between each element of 1hs marray and the rhs scalar. If OP is >>, dataT is a signed type and 1hs marray has a negative value any vacated bits viewed as an unsigned integer must be assigned the value 1, otherwise any vacated bits viewed as an unsigned integer must be assigned the value 0. |
| | Continued on next page |

| Non-member function | Description |
|--|--|
| <pre>marray &operatorOP(marray &lhs, const marray &rhs)</pre> | Available only when: dataT != float && dataT != double && dataT != half. Perform an in-place element-wise OP bit- shift operation between each element of lhs marray and the rhs SYCL marray and returns lhs marray. If OP is >>=, dataT is a signed type and lhs marray has a negative value any vacated bits viewed as an unsigned integer must be assigned the value 1, otherwise any vacated bits viewed as an unsigned integer must be assigned the value 0. |
| | Where OP is: <<=, >>=. |
| <pre>marray &operatorOP(marray &lhs, const dataT &rhs)</pre> | Available only when: dataT != float && dataT != double && dataT != half. Perform an in-place element-wise OP bit- shift operation between each element of lhs marray and the rhs scalar and returns a reference to this SYCL marray. If OP is >>=, dataT is a signed type and lhs marray has a negative value any vacated bits viewed as an unsigned integer must be assigned the value 1, otherwise any vacated bits viewed as an unsigned integer must be assigned the value 0. |
| | Where OP is: <<=, >>=. |
| <pre>marray<bool, numelements=""> operatorOP(const marray& lhs, const marray &rhs)</bool,></pre> | Construct a new instance of the marray class template with dataT = bool and same numElements as lhs marray with each element of the new marray instance the result of an element-wise OP relational operation between each element of lhs marray and each element of the rhs marray. Corresponding element of the marray that is returned must be false if the operation results is a NaN. |
| | Where OP is: ==, !=, <, >, <=, >=. |
| | Continued on next page |

| Non-member function | Description |
|---|---|
| <pre>marray<bool, numelements=""> operatorOP(const marray &</bool,></pre> | Construct a new instance of the marray |
| lhs, const dataT &rhs) | class template with dataT = bool and |
| | same numElements as 1hs marray with |
| | each element of the new marray instance |
| | the result of an element-wise OP relational |
| | operation between each element of 1hs |
| | marray and the rhs scalar. Corresponding |
| | element of the marray that is returned must |
| | be false if the operation results is a NaN. |
| | |
| | Where OP is: ==, !=, <, >, <=, >=. |
| <pre>marray operatorOP(const dataT &lhs, const marray &</pre> | If OP is % available, only when: dataT != |
| rhs) | <pre>float && dataT != double && dataT !=</pre> |
| | half. |
| | Construct a new instance of the SYCL |
| | marray class template with the same tem- |
| | plate parameters as the rhs SYCL marray |
| | with each element of the new SYCL marray |
| | instance the result of an element-wise OP |
| | arithmetic operation between the lhs scalar |
| | and each element of the rhs SYCL marray. |
| | Where OD is: |
| warray anonator OP (const data T like const warray l | Available only when: dataT is float ?? |
| rhe) | dataT l= double && dataT l= half |
| 1113) | Construct a new instance of the SYCI |
| | marray class template with the same tem- |
| | plate parameters as the rhs SYCL marray |
| | with each element of the new SYCL marray |
| | instance the result of an element-wise OP |
| | bitwise operation between the lbs scalar |
| | and each element of the rhs SYCL marray. |
| | |
| | Where OP is: &, , ^. |
| <pre>marray<ret, numelements=""> operatorOP(const dataT &lhs</ret,></pre> | Available only when: dataT != float && |
| , const marray &rhs) | <pre>dataT != double && dataT != half.</pre> |
| | Construct a new instance of the marray |
| | class template with dataT = bool and |
| | same numElements as lhs marray with |
| | each element of the new marray instance |
| | the result of an element-wise OP logical |
| | operation between the 1hs scalar and each |
| | element of the rhs marray. |
| | |
| | Where OP is: &&, . |
| | Continued on next page |

| Non-member function | Description |
|--|---|
| <pre>marray operatorOP(const dataT &lhs, const marray & rhs)</pre> | Construct a new instance of the SYCL marray class template with the same tem- plate parameters as the rhs SYCL marray with each element of the new SYCL marray instance the result of an element-wise OP bitshift operation between the 1hs scalar and each element of the rhs SYCL marray. If OP is >>, dataT is a signed type and this SYCL marray has a negative value any vacated bits viewed as an unsigned integer must be assigned the value 1, otherwise any vacated bits viewed as an unsigned integer must be assigned the value 0. |
| <pre>marray<bool, numelements=""> operatorOP(const dataT & lhs, const marray &rhs)</bool,></pre> | Where OP is: <<, >>. Available only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the marray class template with dataT = bool and same numElements as lhs marray with each element of the new SYCL marray instance the result of an element-wise OP relational operation between the lhs scalar and each element of the rhs marray. Corresponding element of the marray that is returned must be false if the operation results is a NaN. Where OP is: ==, !=, <, >, <=, >=. |
| marray &operator~(const marray &v) | Available only when: dataT != float && dataT != double && dataT != half. Construct a new instance of the SYCL marray class template with the same tem- plate parameters as v marray with each el- ement of the new SYCL marray instance the result of an element-wise OP bitwise opera- tion on each element of v marray. |

| Non-member function | Description |
|--|--|
| <pre>marray<bool, numelements=""> operator!(const marray &v)</bool,></pre> | Construct a new instance of the marray class template with dataT = bool and same numElements as v marray with each element of the new marray instance the result of an element-wise logical ! operation on each element of v marray. |
| | The dataT template parameter of the constructed SYCL marray, RET, varies de- pending on the dataT template parameter of this SYCL marray. For a SYCL marray with dataT of type int8_t or uint8_t RET must be int8_t. For a SYCL marray with dataT of type int16_t, uint16_t or half RET must be int16_t. For a SYCL marray with dataT of type int32_t, uint32_t or float RET must be int32_t. For a SYCL marray with dataT of type int64_t, uint64_t or double RET must be int64_t. |
| | End of table |

Table 4.112: Non-member functions of the marray class template.

4.16.3.2 Aliases

The SYCL programming API provides all permutations of the type alias:

```
using m<type><elems> = marray<<storage-type>, <elems>>
```

where <elems> is 2, 3, 4, 8 and 16, and pairings of <type> and <storage-type> for integral types are char and int8_t, uchar and uint8_t, short and int16_t, ushort and uint16_t, int and int32_t, uint and uint32_t, long and int64_t, ulong and uint64_t, for floating point types are both half, float and double, and for boolean type bool.

For example muint4 is the alias to marray<uint64_t, 4> and mfloat16 is the alias to marray<float, 16>.

4.16.3.3 Memory layout and alignment

The elements of an instance of the marray class template as if stored in std::array<dataT, numElements>.

4.17 Synchronization and atomics

The available features are:

- Accessor classes: Accessor classes specify acquisition and release of buffer and image data structures to provide points at which underlying queue synchronization primitives must be generated.
- Atomic operations: SYCL devices support a restricted subset of C++ atomics and SYCL uses the library syntax from the next C++ specification to make this available.

- Fences: Fence primitives are made available to order loads and stores. They are exposed through the atomic_fence function. Fences can have acquire semantics, release semantics or both.
- Barriers: Barrier primitives are made available to synchronize sets of work-items within individual groups. They are exposed through the group_barrier function.
- Hierarchical parallel dispatch: In the hierarchical parallelism model of describing computations, synchronization within the work-group is made explicit through multiple instances of the parallel_for_work_item function call, rather than through the use of explicit work-group barrier operations.
- Device event: they are used inside SYCL kernel functions to wait for asynchronous operations within a SYCL kernel function to complete.

4.17.1 Barriers and fences

A group barrier or mem-fence provides memory ordering semantics over both the local address space and global address space. All memory operations initiated before the group barrier or mem-fence operation will be completed before any memory operation after the group barrier or mem-fence.

```
1 namespace sycl {
2
3 void atomic_fence(memory_order order, memory_scope scope);
4
5 } // namespace sycl
```

The effects of a call to atomic_fence depend on the value of the order parameter:

- memory_order::relaxed: No effect
- memory_order::acquire: Acquire fence
- memory_order::release: Release fence
- memory_order::acq_rel: Both an acquire fence and a release fence
- memory_order::seq_cst: A sequentially consistent acquire and release fence

A group barrier acts as both an acquire fence and a release fence: all work-items in the group execute a release fence prior to synchronizing at the barrier, and all work-items in the group execute an acquire fence afterwards.

4.17.2 device_event class

The SYCL device_event class encapsulates a single SYCL device event which is available only within SYCL kernel functions and can be used to wait for asynchronous operations within a SYCL kernel function to complete. A SYCL device event can map to a SYCL backend event or just be a host event for a kernel running on the SYCL host device.

All member functions of the device_event class must not throw a SYCL exception.

A synopsis of the SYCL device_event class is provided below. The constructors and member functions of the SYCL device_event class are listed in Table 4.114 and 4.113 respectively.

CHAPTER 4. SYCL PROGRAMMING INTERFACE

324
```
1 namespace sycl {
2 class device_event {
3
4 device_event(__unspecified__);
5
6 public:
7 void wait() noexcept;
8 };
9 } // namespace sycl
```

| Member function | Description |
|--------------------------------|--|
| <pre>void wait()noexcept</pre> | Waits for the asynchronous operation as- |
| | sociated with this SYCL device_event to |
| | complete. |
| | End of table |

Table 4.113: Member functions of the SYCL device_event class.

| Constructor | Description |
|--------------------------------------|---|
| <pre>device_event(unspecified)</pre> | Unspecified implementation defined con- |
| | structor. |
| | End of table |

Table 4.114: Constructors of the device_event class.

4.17.3 Atomic references

The SYCL specification provides atomic operations based on the library syntax from the next C++ specification. The set of supported orderings is specific to a device, but every device is guaranteed to support at least memory_order::relaxed. Since different devices have different capabilities, there is no default ordering in SYCL and the default order used by each instance of sycl::atomic_ref is set by a template argument. If the default order is set to memory_order::relaxed, all memory order arguments default to memory_order::relaxed. If the default order is set to memory_order::acq_rel, memory order arguments default to memory_order::acquire for load operations, memory_order::release for store operations and memory_order::acq_rel for read-modifywrite operations. If the default order is set to memory_order::seq_cst, all memory order arguments default to memory_order::seq_cst.

The SYCL atomic library may map directly to the underlying C++ library in SYCL application code, and must interact safely with the host C++ atomic library when used in host code. The SYCL library must be used in device code to ensure that only the limited subset of functionality is available. SYCL device compilers should give a compilation error on use of the std::atomic and std::atomic_ref classes and functions in device code.

The template parameter addressSpace is permitted to be access::address_space::generic_space, access::address_space::global_space or access::address_space::local_space.

The data type T is permitted to be int, unsigned int, long, unsigned long, long long, unsigned long long, float or double. For floating-point types, the member functions of the atomic_ref class may be emulated, and may use a different floating-point environment to those defined by info::device::single_fp_config and info::device::double_fp_config (i.e. floating-point atomics may use different rounding modes and may have

different exception behavior).

As detailed in Tables 4.116, 4.117, 4.118 and 4.119, not all devices support atomic operations for 64-bit data types. The member functions of the atomic_ref class are required to compile even if the device does not support 64-bit data types, however they are only guaranteed to execute if the device has that support. If a member function is called with a 64-bit data type and the device does not have the necessary support, the SYCL runtime must throw an exception with the errc::feature_not_supported error code.

The atomic types are defined as follows.

```
1 namespace sycl {
 2
   enum class memory_order : /* unspecified */ {
     relaxed, acquire, release, acq_rel, seq_cst
 3
 4 };
 5 inline constexpr memory_order memory_order_relaxed = memory_order::relaxed;
 6 inline constexpr memory_order memory_order_acquire = memory_order::acquire;
 7 inline constexpr memory_order memory_order_release = memory_order::release;
 8 inline constexpr memory_order memory_order_acq_rel = memory_order::acq_rel;
 9 inline constexpr memory_order memory_order_seq_cst = memory_order::seq_cst;
10
11 enum class memory_scope : /* unspecified */ {
12
     work_item, sub_group, work_group, device, system
13 };
14 inline constexpr memory_scope memory_scope_work_item = memory_scope::work_item;
15
   inline constexpr memory_scope memory_scope_sub_group = memory_scope::sub_group;
16 inline constexpr memory_scope memory_scope_work_group = memory_scope::work_group;
17
    inline constexpr memory_scope memory_scope_device = memory_scope::device;
18
    inline constexpr memory_scope memory_scope_system = memory_scope::system;
19
20 // Exposition only
21 template <memory_order ReadModifyWriteOrder>
22 struct memory_order_traits;
23
24 template \Leftrightarrow
25 struct memory_order_traits<memory_order::relaxed> {
26
      static constexpr memory_order read_order = memory_order::relaxed;
27
      static constexpr memory_order write_order = memory_order::relaxed;
28 };
29
30 template <>
31 struct memory_order_traits<memory_order::acq_rel> {
32
     static constexpr memory_order read_order = memory_order::acquire;
33
      static constexpr memory_order write_order = memory_order::release;
34 };
35
36 template <>
37
   struct memory_order_traits<memory_order::seq_cst> {
38
      static constexpr memory_order read_order = memory_order::seq_cst;
39
      static constexpr memory_order write_order = memory_order::seq_cst;
40 };
41
42
    template <typename T, memory_order DefaultOrder, memory_scope DefaultScope, access::address_space</pre>
        Space = access::address_space::generic_space>
43 class atomic_ref {
```

```
44
     public:
45
46
     using value_type = T;
47
      static constexpr size_t required_alignment = /* implementation-defined */;
48
      static constexpr bool is_always_lock_free = /* implementation-defined */;
49
      static constexpr memory_order default_read_order = memory_order_traits<DefaultOrder>::read_order
50
      static constexpr memory_order default_write_order = memory_order_traits<DefaultOrder>::
          write_order;
51
      static constexpr memory_order default_read_modify_write_order = DefaultOrder;
52
      static constexpr memory_scope default_scope = DefaultScope;
53
54
      bool is_lock_free() const noexcept;
55
56
      explicit atomic_ref(T&);
57
      atomic_ref(const atomic_ref&) noexcept;
58
      atomic_ref& operator=(const atomic_ref&) = delete;
59
60
      void store(T operand,
61
        memory_order order = default_write_order,
       memory_scope scope = default_scope) const noexcept;
62
63
64
      T operator=(T desired) const noexcept;
65
66
      T load(memory_order order = default_read_order,
67
        memory_scope scope = default_scope) const noexcept;
68
69
      operator T() const noexcept;
70
71
     T exchange(T operand,
72
        memory_order order = default_read_modify_write_order,
73
        memory_scope scope = default_scope) const noexcept;
74
75
      bool compare_exchange_weak(T &expected, T desired,
76
        memory_order success,
77
        memory_order failure,
78
        memory_scope scope = default_scope) const noexcept;
79
80
      bool compare_exchange_weak(T &expected, T desired,
81
        memory_order order = default_read_modify_write_order,
82
        memory_scope scope = default_scope) const noexcept;
83
84
     bool compare_exchange_strong(T &expected, T desired,
85
        memory_order success,
86
        memory_order failure,
87
        memory_scope scope = default_scope) const noexcept;
88
      bool compare_exchange_strong(T &expected, T desired,
89
90
        memory_order order = default_read_modify_write_order,
91
        memory_scope scope = default_scope) const noexcept;
92 };
93
94
   // Partial specialization for integral types
95 template <memory_order DefaultOrder, memory_scope DefaultScope, access::address_space Space =
        access::address_space::generic_space>
```

```
96 class atomic_ref<Integral, DefaultOrder, DefaultScope, Space> {
97
98
       /* All other members from atomic_ref<T> are available */
99
100
      using difference_type = value_type;
101
102
       Integral fetch_add(Integral operand,
103
         memory_order order = default_read_modify_write_order,
104
         memory_scope scope = default_scope) const noexcept;
105
106
       Integral fetch_sub(Integral operand,
107
         memory_order order = default_read_modify_write_order,
108
         memory_scope scope = default_scope) const noexcept;
109
110
       Integral fetch_and(Integral operand,
111
         memory_order order = default_read_modify_write_order,
112
         memory_scope scope = default_scope) const noexcept;
113
114
       Integral fetch_or(Integral operand,
115
         memory_order order = default_read_modify_write_order,
116
         memory_scope scope = default_scope) const noexcept;
117
       Integral fetch_min(Integral operand,
118
119
         memory_order order = default_read_modify_write_order,
120
         memory_scope scope = default_scope) const noexcept;
121
122
       Integral fetch_max(Integral operand,
123
         memory_order order = default_read_modify_write_order,
124
         memory_scope scope = default_scope) const noexcept;
125
126
       Integral operator++(int) const noexcept;
127
       Integral operator--(int) const noexcept;
128
       Integral operator++() const noexcept;
129
       Integral operator--() const noexcept;
130
       Integral operator+=(Integral) const noexcept;
131
       Integral operator-=(Integral) const noexcept;
132
       Integral operator&=(Integral) const noexcept;
133
       Integral operator|=(Integral) const noexcept;
134
       Integral operator^=(Integral) const noexcept;
135
136 };
137
138 // Partial specialization for floating-point types
139
    template <memory_order DefaultOrder, memory_scope DefaultScope, access::address_space Space =</pre>
         access::address_space::generic_space>
140
    class atomic_ref<Floating, DefaultOrder, DefaultScope, Space> {
141
       /* All other members from atomic_ref<T> are available */
142
143
144
       using difference_type = value_type;
145
146
       Floating fetch_add(Floating operand,
147
         memory_order order = default_read_modify_write_order,
148
         memory_scope scope = default_scope) const noexcept;
149
```

```
150
       Floating fetch_sub(Floating operand,
151
         memory_order order = default_read_modify_write_order,
152
         memory_scope scope = default_scope) const noexcept;
153
154
      Floating fetch_min(Floating operand,
155
         memory_order order = default_read_modify_write_order,
156
         memory_scope scope = default_scope) const noexcept;
157
158
       Floating fetch_max(Floating operand,
159
         memory_order order = default_read_modify_write_order,
160
         memory_scope scope = default_scope) const noexcept;
161
162
       Floating operator+=(Floating) const noexcept;
163
       Floating operator-=(Floating) const noexcept;
164
165 };
166
167
    // Partial specialization for pointers
168
    template <typename T, memory_order DefaultOrder, memory_scope DefaultScope, access::address_space
         Space = access::address_space::generic_space>
169
     class atomic_ref<T*, DefaultOrder, DefaultScope, Space> {
170
171
      using value_type = T*;
172
      using difference_type = ptrdiff_t;
       static constexpr size_t required_alignment = /* implementation-defined */;
173
       static constexpr bool is_always_lock_free = /* implementation-defined */;
174
175
       static constexpr memory_order default_read_order = memory_order_traits<DefaultOrder>::read_order
176
       static constexpr memory_order default_write_order = memory_order_traits<DefaultOrder>::
           write_order:
177
       static constexpr memory_order default_read_modify_write_order = DefaultOrder;
178
       static constexpr memory_scope default_scope = DefaultScope;
179
180
       bool is_lock_free() const noexcept;
181
182
       explicit atomic_ref(T*&);
183
       atomic_ref(const atomic_ref&) noexcept;
184
       atomic_ref& operator=(const atomic_ref&) = delete;
185
186
       void store(T* operand,
         memory_order order = default_write_order,
187
188
         memory_scope scope = default_scope) const noexcept;
189
190
       T* operator=(T* desired) const noexcept;
191
192
       T* load(memory_order order = default_read_order,
193
         memory_scope scope = default_scope) const noexcept;
194
195
       operator T*() const noexcept;
196
197
       T* exchange(T* operand,
198
         memory_order order = default_read_modify_write_order,
199
         memory_scope scope = default_scope) const noexcept;
200
201
       bool compare_exchange_weak(T* &expected, T* desired,
```

```
202
         memory_order success,
203
         memory order failure.
204
         memory_scope scope = default_scope) const noexcept;
205
206
       bool compare_exchange_weak(T* &expected, T* desired,
         memory_order order = default_read_modify_write_order,
207
208
         memory_scope scope = default_scope) const noexcept;
209
       bool compare_exchange_strong(T* &expected, T* desired,
210
211
        memory_order success,
212
        memory_order failure,
213
        memory_scope scope = default_scope) const noexcept;
214
215
       bool compare_exchange_strong(T* &expected, T* desired,
         memory_order order = default_read_modify_write_order,
216
217
         memory_scope scope = default_scope) const noexcept;
218
      T* fetch_add(difference_type,
219
220
         memory_order order = default_read_modify_write_order,
221
         memory_scope scope = default_scope) const noexcept;
222
223
      T* fetch_sub(difference_type,
224
         memory_order order = default_read_modify_write_order,
225
        memory_scope scope = default_scope) const noexcept;
226
227
      T* operator++(int) const noexcept;
228
      T* operator--(int) const noexcept;
229
      T* operator++() const noexcept;
230
      T* operator--() const noexcept;
231
      T* operator+=(difference_type) const noexcept;
232
       T* operator-=(difference_type) const noexcept;
233
234 };
235
236 } // namespace sycl
```

The constructors and member functions for instances of the SYCL atomic_ref class using any compatible type are listed in Tables 4.115 and 4.116 respectively. Additional member functions for integral, floating-point and pointer types are listed in Tables 4.117, 4.118 and 4.119 respectively.

The static member required_alignment describes the minimum required alignment in bytes of an object that can be referenced by an atomic_ref<T>, which must be at least alignof(T).

The static member is_always_lock_free is true if all atomic operations for type T are always lock-free. A SYCL implementation is not guaranteed to support atomic operations that are not lock-free.

The static members default_read_order, default_write_order and default_read_modify_write_order reflect the default memory order values for each type of atomic operation, consistent with the DefaultOrder template.

The atomic operations and member functions behave as described in the C++ specification, barring the restrictions discussed above. Note that care must be taken when using atomics to implement synchronization routines due to the lack of forward progress guarantees between work-items in SYCL. No work-item may be dependent on another work-item to make progress if the code is to be portable.

| Constructor | Description |
|-----------------------------------|-------------------------------------|
| <pre>atomic_ref(T& ref)</pre> | Constructs an instance of SYCL |
| | atomic_ref which is associated with |
| | the reference ref. |
| | End of table |

Table 4.115: Constructors of the SYCL atomic_ref class template.

| Member function | Description |
|--|--|
| <pre>bool is_lock_free()const</pre> | Return true if the atomic operations pro- |
| | vided by this atomic_ref are lock-free. |
| <pre>void store(T operand,</pre> | Atomically stores operand to the object ref- |
| <pre>memory_order order = default_write_order,</pre> | erenced by this atomic_ref. The mem- |
| <pre>memory_scope scope = default_scope)const</pre> | ory order of this atomic operation must be |
| | <pre>memory_order::relaxed, memory_order::</pre> |
| | release or memory_order::seq_cst. This |
| | function is only supported for 64-bit |
| | data types on devices that have aspect:: |
| | int64_base_atomics. |
| T operator=(T desired)const | Equivalent to store(desired). Returns |
| | desired. |
| T load(| Atomically loads the value of the object ref- |
| <pre>memory_order order = default_read_order</pre> | erenced by this atomic_ref. The mem- |
| <pre>memory_odrer scope = default_scope)const</pre> | ory order of this atomic operation must be |
| | <pre>memory_order::relaxed, memory_order::</pre> |
| | acquire, or memory_order::seq_cst. This |
| | function is only supported for 64-bit |
| | data types on devices that have aspect:: |
| | int64_base_atomics. |
| operator T()const | Equivalent to load(). |
| T exchange(T operand, | Atomically replaces the value of the object |
| <pre>memory_order order =</pre> | referenced by this atomic_ref with value |
| <pre>default_read_modify_write_order,</pre> | operand and returns the original value of the |
| <pre>memory_scope scope = default_scope)const</pre> | referenced object. This function is only sup- |
| | ported for 64-bit data types on devices that |
| | have aspect::int64_base_atomics. |
| | Continued on next page |

Table 4.116: Member functions available on any object of type atomic_ref<T>.

| Member function | Description |
|---|--|
| <pre>bool compare_exchange_weak(T &expected, T desired,</pre> | Atomically compares the value of the object |
| <pre>memory_order success,</pre> | referenced by this atomic_ref against the |
| memory_order failure, | value of expected. If the values are equal, |
| <pre>memory_scope scope = default_scope)const</pre> | attempts to replace the value of the refer- |
| | enced object with the value of desired; oth- |
| | erwise assigns the original value of the ref- |
| | erenced object to expected. |
| | Returns true if the comparison operation |
| | and replacement operation were successful. |
| | The failure memory order of this atomic |
| | operation must be memory_order::relaxed |
| | <pre>, memory_order::acquire or memory_order</pre> |
| | ::seq_cst. |
| | This function is only supported for 64-bit |
| | data types on devices that have aspect:: |
| | int64_base_atomics. |
| <pre>bool compare_exchange_weak(T &expected, T desired,</pre> | Equivalent to compare_exchange_weak |
| <pre>memory_order order =</pre> | (expected, desired, order, order, |
| <pre>default_read_modify_write_order,</pre> | scope). |
| <pre>memory_scope scope = default_scope)const</pre> | |
| <pre>bool compare_exchange_strong(T &expected, T desired,</pre> | Atomically compares the value of the object |
| memory_order success, | referenced by this atomic_ref against the |
| memory_order failure, | value of expected. If the values are equal, |
| <pre>memory_scope scope = default_scope)const</pre> | replaces the value of the referenced object |
| | with the value of desired; otherwise assigns |
| | the original value of the referenced object to |
| | expected. |
| | Returns true if the comparison opera- |
| | tion was successful. The failure mem- |
| | ory order of this atomic operation must |
| | be memory_order::relaxed, memory_order |
| | ::acquire or memory_order::seq_cst. |
| | This function is only supported for 64-bit |
| | data types on devices that have aspect:: |
| | int64_base_atomics. |
| <pre>bool compare_exchange_strong(T &expected, T desired,</pre> | Equivalent to compare_exchange_strong |
| <pre>memory_order order =</pre> | (expected, desired, order, order, |
| <pre>default_read_modify_write_order)const</pre> | scope). |
| | End of table |

Table 4.116: Member functions available on any object of type atomic_ref<T>.

| Member function | Description |
|---|---|
| T fetch_add(T operand, | Atomically adds operand to the value of the |
| <pre>memory_order order =</pre> | object referenced by this atomic_ref and as- |
| <pre>default_read_modify_write_order,</pre> | signs the result to the value of the referenced |
| <pre>memory_scope scope = default_scope)const</pre> | object. Returns the original value of the ref- |
| | erenced object. This function is only sup- |
| | ported for 64-bit data types on devices that |
| | have aspect::int64_base_atomics. |
| T operator+=(T operand)const | Equivalent to fetch_add(operand). |
| T operator++(int)const | Equivalent to fetch_add(1). |
| T operator++()const | Equivalent to fetch_add(1)+ 1. |
| T fetch_sub(T operand, | Atomically subtracts operand from the value |
| <pre>memory_order order =</pre> | of the object referenced by this atomic_ref |
| <pre>default_read_modify_write_order,</pre> | and assigns the result to the value of the ref- |
| <pre>memory_scope scope = default_scope)const</pre> | erenced object. Returns the original value of |
| | the referenced object. This function is only |
| | supported for 64-bit data types on devices |
| | that have aspect::int64_base_atomics. |
| T operator-=(T operand)const | Equivalent to fetch_sub(operand). |
| T operator(int)const | Equivalent to fetch_sub(1). |
| T operator()const | Equivalent to fetch_sub(1)- 1. |
| T fetch_and(T operand, | Atomically performs a bitwise AND be- |
| memory_order order = | tween operand and the value of the object |
| <pre>default_read_modify_write_order,</pre> | referenced by this atomic_ref, and assigns |
| <pre>memory_scope scope = default_scope)const</pre> | the result to the value of the referenced ob- |
| | ject. Returns the original value of the ref- |
| | erenced object. This function is only sup- |
| | ported for 64-bit data types on devices that |
| | have aspect::int64_extended_atomics. |
| T operator&=(T operand)const | Equivalent to fetch_and(operand). |
| T fetch_or(T operand, | Atomically performs a bitwise OR between |
| <pre>memory_order order =</pre> | operand and the value of the object refer- |
| <pre>default_read_modify_write_order,</pre> | enced by this atomic_ref, and assigns the |
| <pre>memory_scope scope = default_scope)const</pre> | result to the value of the referenced ob- |
| | ject. Returns the original value of the ref- |
| | erenced object. This function is only sup- |
| | ported for 64-bit data types on devices that |
| | have aspect::int64_extended_atomics. |
| T operator = (T operand) const | Equivalent to fetch_or(operand). |
| T fetch_xor(T operand, | Atomically performs a bitwise XOR be- |
| memory_order order = | tween operand and the value of the object |
| <pre>default_read_modify_write_order,</pre> | referenced by this atomic_ref, and assigns |
| <pre>memory_scope scope = default_scope)const</pre> | the result to the value of the referenced ob- |
| | ject. Returns the original value of the ref- |
| | erenced object. This function is only sup- |
| | ported for 64-bit data types on devices that |
| | have aspect::int64_extended_atomics. |
| T operator [^] =(T operand)const | Equivalent to fetch_xor(operand). |
| | Continued on next page |

Table 4.117: Additional member functions available on an object of type atomic_ref<T> for integral T.

| Member function | Description |
|---|--|
| T fetch_min(T operand, | Atomically computes the minimum of |
| <pre>memory_order order =</pre> | operand and the value of the object refer- |
| <pre>default_read_modify_write_order,</pre> | enced by this atomic_ref, and assigns the |
| <pre>memory_scope scope = default_scope)const</pre> | result to the value of the referenced ob- |
| | ject. Returns the original value of the ref- |
| | erenced object. This function is only sup- |
| | ported for 64-bit data types on devices that |
| | have aspect::int64_extended_atomics. |
| T fetch_max(T operand, | Atomically computes the maximum of |
| <pre>memory_order order =</pre> | operand and the value of the object refer- |
| <pre>default_read_modify_write_order,</pre> | enced by this atomic_ref, and assigns the |
| <pre>memory_scope scope = default_scope)const</pre> | result to the value of the referenced ob- |
| | ject. Returns the original value of the ref- |
| | erenced object. This function is only sup- |
| | ported for 64-bit data types on devices that |
| | have aspect::int64_extended_atomics. |
| | End of table |

Table 4.117: Additional member functions available on an object of type atomic_ref<T> for integral T.

| Member function | Description |
|---|---|
| T fetch_add(T operand, | Atomically adds operand to the value of the |
| <pre>memory_order order =</pre> | object referenced by this atomic_ref and as- |
| <pre>default_read_modify_write_order,</pre> | signs the result to the value of the referenced |
| <pre>memory_scope scope = default_scope)const</pre> | object. Returns the original value of the ref- |
| | erenced object. This function is only sup- |
| | ported for 64-bit data types on devices that |
| | have aspect::int64_base_atomics. |
| T operator+=(T operand)const | Equivalent to fetch_add(operand). |
| T fetch_sub(T operand, | Atomically subtracts operand from the value |
| <pre>memory_order order =</pre> | of the object referenced by this atomic_ref |
| <pre>default_read_modify_write_order,</pre> | and assigns the result to the value of the ref- |
| <pre>memory_scope scope = default_scope)const</pre> | erenced object. Returns the original value of |
| | the referenced object. This function is only |
| | supported for 64-bit data types on devices |
| | that have aspect::int64_base_atomics. |
| T operator-=(T operand)const | Equivalent to fetch_sub(operand). |
| T fetch_min(T operand, | Atomically computes the minimum of |
| <pre>memory_order order =</pre> | operand and the value of the object refer- |
| <pre>default_read_modify_write_order,</pre> | enced by this atomic_ref, and assigns the |
| <pre>memory_scope scope = default_scope)const</pre> | result to the value of the referenced ob- |
| | ject. Returns the original value of the ref- |
| | erenced object. This function is only sup- |
| | ported for 64-bit data types on devices that |
| | have aspect::int64_extended_atomics. |
| | Continued on next page |

Table 4.118: Additional member functions available on an object of type atomic_ref<T> for floating-point T.

| Member function | Description |
|---|--|
| T fetch_max(T operand, | Atomically computes the maximum of |
| <pre>memory_order order =</pre> | operand and the value of the object refer- |
| <pre>default_read_modify_write_order,</pre> | enced by this atomic_ref, and assigns the |
| <pre>memory_scope scope = default_scope)const</pre> | result to the value of the referenced ob- |
| | ject. Returns the original value of the ref- |
| | erenced object. This function is only sup- |
| | ported for 64-bit data types on devices that |
| | have aspect::int64_extended_atomics. |
| | End of table |

Table 4.118: Additional member functions available on an object of type atomic_ref<T> for floating-point T.

| Member function | Description |
|---|---|
| T* fetch_add(ptrdiff_t operand, | Atomically adds operand to the value of the |
| <pre>memory_order order =</pre> | object referenced by this atomic_ref and |
| <pre>default_read_modify_write_order,</pre> | assigns the result to the value of the refer- |
| <pre>memory_scope scope = default_scope)const</pre> | enced object. Returns the original value of |
| | the referenced object. This function is only |
| | supported for 64-bit pointers on devices that |
| | have aspect::int64_base_atomics. |
| T* operator+=(ptrdiff_t operand)const | Equivalent to fetch_add(operand). |
| T* operator++(int)const | Equivalent to fetch_add(1). |
| T* operator++()const | Equivalent to fetch_add(1)+ 1. |
| T* fetch_sub(ptrdiff_t operand, | Atomically subtracts operand from the value |
| <pre>memory_order order =</pre> | of the object referenced by this atomic_ref |
| <pre>default_read_modify_write_order,</pre> | and assigns the result to the value of the ref- |
| <pre>memory_scope scope = default_scope)const</pre> | erenced object. Returns the original value of |
| | the referenced object. This function is only |
| | supported for 64-bit pointers on devices that |
| | have aspect::int64_base_atomics. |
| T* operator-=(ptrdiff_t operand)const | Equivalent to fetch_sub(operand). |
| T* operator(int)const | Equivalent to fetch_sub(1). |
| T* operator()const | Equivalent to fetch_sub(1)- 1. |
| | End of table |

Table 4.119: Additional member functions available on an object of type atomic_ref<T*>.

4.17.4 Atomic types (deprecated)

The atomic types and operations on atomic types provided by SYCL 1.2.1 are deprecated in SYCL 2020, and will be removed in a future version of SYCL. The types and operations are made available in the cl::sycl:: namespace for backwards compatibility.

The constructors and member functions for the cl::sycl::atomic class are listed in Tables 4.120 and 4.121 respectively.

```
1 namespace cl {
```

```
2 namespace sycl {
   /* Deprecated in SYCL 2020 */
 3
 4
    enum class memory_order : int {
 5
      relaxed
 6 };
 7
 8
   /* Deprecated in SYCL 2020 */
    template <typename T, access::address_space addressSpace =</pre>
 9
      access::address_space::global_space>
10
11 class atomic {
12
     public:
      template <typename pointerT, access::decorated IsDecorated>
13
      atomic(multi_ptr<pointerT, addressSpace, IsDecorated> ptr);
14
15
      void store(T operand, memory_order memoryOrder =
16
17
        memory_order::relaxed);
18
19
      T load(memory_order memoryOrder = memory_order::relaxed) const;
20
21
      T exchange(T operand, memory_order memoryOrder =
22
       memory_order::relaxed);
23
24
      /* Available only when: T != float */
25
      bool compare_exchange_strong(T &expected, T desired,
26
        memory_order successMemoryOrder = memory_order::relaxed,
27
        memory_order failMemoryOrder = memory_order::relaxed);
28
29
      /* Available only when: T != float */
30
      T fetch_add(T operand, memory_order memoryOrder =
31
        memory_order::relaxed);
32
33
      /* Available only when: T != float */
34
      T fetch_sub(T operand, memory_order memoryOrder =
35
        memory_order::relaxed);
36
37
      /* Available only when: T != float */
38
      T fetch_and(T operand, memory_order memoryOrder =
39
        memory_order::relaxed);
40
      /* Available only when: T != float */
41
42
      T fetch_or(T operand, memory_order memoryOrder =
43
        memory_order::relaxed);
44
45
      /* Available only when: T != float */
      T fetch_xor(T operand, memory_order memoryOrder =
46
47
        memory_order::relaxed);
48
49
      /* Available only when: T != float */
50
      T fetch_min(T operand, memory_order memoryOrder =
51
        memory_order::relaxed);
52
53
      /* Available only when: T != float */
      T fetch_max(T operand, memory_order memoryOrder =
54
55
        memory_order::relaxed);
56 };
```

```
57 } // namespace sycl
58 } // namespace cl
```

The global functions are as follows and described in Table 4.122.

```
1 namespace cl {
2 namespace sycl {
3 /* Deprecated in SYCL 2020 */
4 template <typename T, access::address_space addressSpace>
5 void atomic_store(atomic<T, addressSpace> object, T operand, memory_order memoryOrder =
6
     memory_order::relaxed);
7
8
   /* Deprecated in SYCL 2020 */
9
   template <typename T, access::address_space addressSpace>
10 T atomic_load(atomic<T, addressSpace> object, memory_order memoryOrder =
     memory_order::relaxed);
11
12
13 /* Deprecated in SYCL 2020 */
14 template <typename T, access::address_space addressSpace>
15 T atomic_exchange(atomic<T, addressSpace> object, T operand, memory_order memoryOrder =
16
     memory_order::relaxed);
17
18 /* Deprecated in SYCL 2020 */
19
   template <typename T, access::address_space addressSpace>
20
   bool atomic_compare_exchange_strong(atomic<T, addressSpace> object, T &expected, T desired,
21
       memory_order successMemoryOrder = memory_order::relaxed,
22
       memory_order failMemoryOrder = memory_order::relaxed);
23
24
   /* Deprecated in SYCL 2020 */
25
   template <typename T, access::address_space addressSpace>
26
   T atomic_fetch_add(atomic<T, addressSpace> object, T operand, memory_order memoryOrder =
27
       memory_order::relaxed);
28
29 /* Deprecated in SYCL 2020 */
30 template <typename T, access::address_space addressSpace>
31 T atomic_fetch_sub(atomic<T, addressSpace> object, T operand, memory_order memoryOrder =
32
     memory_order::relaxed);
33
34 /* Deprecated in SYCL 2020 */
35
   template <typename T, access::address_space addressSpace>
36 T atomic_fetch_and(atomic<T, addressSpace> object, T operand, memory_order memoryOrder =
37
     memory_order::relaxed);
38
39 /* Deprecated in SYCL 2020 */
40 template <typename T, access::address_space addressSpace>
41 T atomic_fetch_or(atomic<T, addressSpace> object, T operand, memory_order memoryOrder =
42
    memory_order::relaxed);
43
44 /* Deprecated in SYCL 2020 */
45 template <typename T, access::address_space addressSpace>
46 T atomic_fetch_xor(atomic<T, addressSpace> object, T operand, memory_order memoryOrder =
47
     memory_order::relaxed);
48
49 /* Deprecated in SYCL 2020 */
```

50 template <typename T, access::address_space addressSpace>

```
51 T atomic_fetch_min(atomic<T, addressSpace> object, T operand, memory_order memoryOrder =
```

```
52 memory_order::relaxed);
```

```
53
```

```
54 /* Deprecated in SYCL 2020 */
```

55 template <typename T, access::address_space addressSpace>

- 56 T atomic_fetch_max(atomic<T, addressSpace> object, T operand, memory_order memoryOrder =
- 57 memory_order::relaxed);

58 } // namespace sycl

```
59 } // namespace cl
```

| Constructor | Description |
|---|--|
| <pre>template <typename pointert=""></typename></pre> | Deprecated in SYCL 2020. |
| <pre>atomic(multi_ptr<pointert, addressspace=""> ptr)</pointert,></pre> | Permitted data types for pointerT are any valid scalar data type which is the same size in bytes as T. Constructs an instance of SYCL atomic which is associated with the pointer ptr, converted to a pointer of data type T. |
| | End of table |

Table 4.120: Constructors of the cl::sycl::atomic class template.

| Member function | Description |
|---|--|
| <pre>void store(T operand, memory_order memoryOrder =</pre> | Deprecated in SYCL 2020. |
| <pre>memory_order::relaxed)</pre> | Atomically stores the value operand at the |
| | address of the multi_ptr associated with |
| | this SYCL atomic. The memory order of |
| | this atomic operation must be memory_order |
| | ::relaxed. This function is only supported |
| | for 64-bit data types on devices that have |
| | <pre>aspect::int64_base_atomics.</pre> |
| T load(memory_order memoryOrder = | Deprecated in SYCL 2020. |
| <pre>memory_order::relaxed)const</pre> | Atomically loads the value at the address of |
| | the multi_ptr associated with this SYCL |
| | atomic. Returns the value at the address |
| | of the multi_ptr associated with this SYCL |
| | atomic before the call. The memory order of |
| | this atomic operation must be memory_order |
| | ::relaxed. This function is only supported |
| | for 64-bit data types on devices that have |
| | <pre>aspect::int64_base_atomics.</pre> |
| | Continued on next page |

| Member function | Description |
|---|---|
| T exchange(T operand, memory_order memoryOrder = | Deprecated in SYCL 2020. |
| <pre>memory_order::relaxed)</pre> | Atomically replaces the value at the address |
| | of the multi_ptr associated with this SYCL |
| | atomic with value operand and returns the |
| | value at the address of the multi_ptr as- |
| | sociated with this SYCL atomic before the |
| | call. The memory order of this atomic |
| | operation must be memory_order::relaxed |
| | . This function is only supported for 64- |
| | bit data types on devices that have aspect |
| | ::int64_base_atomics. |
| <pre>bool compare_exchange_strong(T &expected, T desired,</pre> | Deprecated in SYCL 2020. |
| <pre>memory_order successMemoryOrder =</pre> | Available only when: T != float. |
| <pre>memory_order::relaxed,</pre> | Atomically compares the value at the ad- |
| memory_order failMemoryOrder = | dress of the multi_ptr associated with |
| memory_order::relaxed) | this SYCL atomic against the value of |
| | value at address of the multing the associ |
| | ated with this SVCL atomic with the value |
| | of desired: otherwise assigns the origi- |
| | nal value at the address of the multi ntr |
| | associated with this SYCL atomic to |
| | expected. Returns true if the compari- |
| | son operation was successful. The mem- |
| | ory order of this atomic operation must |
| | be memory_order::relaxed for both success |
| | and fail. This function is only supported |
| | for 64-bit data types on devices that have |
| | <pre>aspect::int64_base_atomics.</pre> |
| T fetch_add(T operand, memory_order memoryOrder = | Deprecated in SYCL 2020. |
| <pre>memory_order::relaxed)</pre> | Available only when: T != float. |
| | Atomically adds the value operand to the |
| | value at the address of the multi_ptr as- |
| | sociated with this SYCL atomic and as- |
| | signs the result to the value at the address |
| | of the multi_ptr associated with this SYCL |
| | atomic. Returns the value at the address |
| | of the multi_ptr associated with this SYCL |
| | this atomic operation must be memory order of |
| | uns atomic operation must be memory_order |
| | for 64-bit data types on devices that have |
| | aspect: int64 base atomics |
| | Continued on next page |

| T fetch_sub(T operand, memory_order memoryOrder = Deprecated in SYCL 2020. | |
|--|-------------|
| | |
| memory_order::relaxed) Available only when: T != float. | |
| Atomically subtracts the value operand | d to |
| the value at the address of the multi_ | ptr |
| associated with this SYCL atomic and | as- |
| signs the result to the value at the addr | ress |
| of the multi_ptr associated with this SY | ′CL |
| atomic. Returns the value at the add | ress |
| of the multi_ptr associated with this SY | 'CL |
| atomic before the call. The memory orde | er of |
| this atomic operation must be memory_or | der |
| ::relaxed. This function is only support | rted |
| for 64-bit data types on devices that h | ave |
| aspect::int64_base_atomics. | |
| T fetch_and(T operand, memory_order memoryOrder = Deprecated in SYCL 2020. | |
| memory_order::relaxed) Available only when: T != float. | 1. |
| Atomically performs a bitwise AND | be- |
| the oddress of the multi- new associated w | |
| the address of the multi_ptr associated v | |
| the value at the address of the multi-inter | |
| sociated with this SVCL atomic. Retu | as- irne |
| the value at the address of the multi | ntr |
| associated with this SYCL atomic be | fore |
| the call. The memory order of this ato | mic |
| operation must be memory order::rela | xed |
| . This function is only supported for | 64- |
| bit data types on devices that have asp | ect |
| ::int64_extended_atomics. | |
| T fetch_or(T operand, memory_order memoryOrder = Deprecated in SYCL 2020. | |
| memory_order::relaxed) Available only when: T != float. | |
| Atomically performs a bitwise OR betw | veen |
| the value operand and the value at the | ad- |
| dress of the multi_ptr associated with | this |
| SYCL atomic and assigns the result to | the |
| value at the address of the multi_ptr | as- |
| sociated with this SYCL atomic. Retu | ırns |
| the value at the address of the multi_ | ptr |
| associated with this SYCL atomic bet | tore |
| the call. The memory order of this atom | mic |
| operation must be memory_order::rela | xed |
| . This function is only supported for | 64- |
| bit data types on devices that have asp | ect |
| ::Into4_extended_atomics. | 0.000 |

| Member function | Description |
|---|---|
| T fetch_xor(T operand, memory_order memoryOrder = | Deprecated in SYCL 2020. |
| <pre>memory_order::relaxed)</pre> | Available only when: T != float. |
| | Atomically performs a bitwise XOR be- |
| | tween the value operand and the value at |
| | the address of the multi_ptr associated with |
| | this SYCL atomic and assigns the result to |
| | the value at the address of the multi_ptr as- |
| | sociated with this SYCL atomic. Returns |
| | the value at the address of the multi_ptr |
| | associated with this SYCL atomic before |
| | the call. The memory order of this atomic |
| | operation must be memory_order::relaxed |
| | . This function is only supported for 64- |
| | bit data types on devices that have aspect |
| | ::int64_extended_atomics. |
| T fetch_min(T operand, memory_order memoryOrder = | Deprecated in SYCL 2020. |
| <pre>memory_order::relaxed)</pre> | Atomically computes the minimum of the |
| | value operand and the value at the ad- |
| | dress of the multi_ptr associated with this |
| | sich atomic and assigns the result to the |
| | sociated with this SVCI atomic Returns |
| | the value at the address of the multi ntr |
| | associated with this SVCL atomic before |
| | the call The memory order of this atomic |
| | operation must be memory order::relaxed |
| | . This function is only supported for 64- |
| | bit data types on devices that have aspect |
| | ::int64_extended_atomics. |
| T fetch_max(T operand, memory_order memoryOrder = | Deprecated in SYCL 2020. |
| <pre>memory_order::relaxed)</pre> | Available only when: T != float. |
| | Atomically computes the maximum of the |
| | value operand and the value at the ad- |
| | dress of the multi_ptr associated with this |
| | SYCL atomic and assigns the result to the |
| | value at the address of the multi_ptr as- |
| | sociated with this SYCL atomic. Returns |
| | the value at the address of the multi_ptr |
| | associated with this SYCL atomic before |
| | the call. The memory order of this atomic |
| | operation must be memory_order::relaxed |
| | . This function is only supported for 64- |
| | bit data types on devices that have aspect |
| | ::int64_extended_atomics. |
| | End of table |

| Functions | Description |
|--|---|
| <pre>template <typename access::address_space<="" pre="" t,=""></typename></pre> | Deprecated in SYCL 2020. |
| addressSpace> | Equivalent to calling object.load(|
| T atomic_load(atomic <t, addressspace=""> <pre>object,</pre></t,> | memoryOrder). |
| <pre>memory_order memoryOrder = memory_order::relaxed</pre> | |
|) | |
| <pre>template <typename access::address_space<="" pre="" t,=""></typename></pre> | Deprecated in SYCL 2020. |
| addressSpace> | Equivalent to calling object.store(|
| <pre>void atomic_store(atomic<t, addressspace=""> object, T</t,></pre> | operand, memoryOrder). |
| operand, | |
| <pre>memory_order memoryOrder = memory_order::relaxed</pre> | |
|) | |
| <pre>template <typename access::address_space<="" pre="" t,=""></typename></pre> | Deprecated in SYCL 2020. |
| addressSpace> | Equivalent to calling object.exchange(|
| <pre>T atomic_exchange(atomic<t, addressspace=""> object, T</t,></pre> | operand, memoryOrder). |
| operand, | |
| <pre>memory_order memoryOrder = memory_order::relaxed</pre> | |
|) | |
| <pre>template <typename access::address_space<="" pre="" t,=""></typename></pre> | Deprecated in SYCL 2020. |
| addressSpace> | Equivalent to calling object. |
| <pre>bool atomic_compare_exchange_strong(</pre> | <pre>compare_exchange_strong(expected</pre> |
| atomic <t, addressspace=""> object, T &expected, T</t,> | <pre>, desired, successMemoryOrder,</pre> |
| desired, | failMemoryOrders). |
| <pre>memory_order successMemoryOrder =</pre> | |
| <pre>memory_order::relaxed</pre> | |
| <pre>memory_order failMemoryOrder =</pre> | |
| <pre>memory_order::relaxed)</pre> | |
| <pre>template <typename access::address_space<="" pre="" t,=""></typename></pre> | Deprecated in SYCL 2020. |
| addressSpace> | Equivalent to calling object.fetch_add(|
| T atomic_fetch_add(atomic <t, addressspace=""> object, T</t,> | operand, memoryOrder). |
| operand, | |
| <pre>memory_order memoryOrder = memory_order::relaxed</pre> | |
|) | |
| <pre>template <typename access::address_space<="" pre="" t,=""></typename></pre> | Deprecated in SYCL 2020. |
| addressSpace> | Equivalent to calling object.fetch_sub(|
| <pre>T atomic_fetch_sub(atomic<t, addressspace=""> object, T</t,></pre> | operand, memoryOrder). |
| operand, | |
| <pre>memory_order memoryOrder = memory_order::relaxed</pre> | |
|) | |
| <pre>template <typename access::address_space<="" pre="" t,=""></typename></pre> | Deprecated in SYCL 2020. |
| addressSpace> | Equivalent to calling object.fetch_add(|
| <pre>T atomic_fetch_and(atomic<t> operand, T object,</t></pre> | operand, memoryOrder). |
| <pre>memory_order memoryOrder = memory_order::relaxed</pre> | |
|) | |
| | Continued on next page |

Table 4.122: Global functions available on atomic types.

| Functions | Description |
|--|--|
| <pre>template <typename access::address_space<="" pre="" t,=""></typename></pre> | Deprecated in SYCL 2020. |
| addressSpace> | Equivalent to calling <pre>object.fetch_or(</pre> |
| <pre>T atomic_fetch_or(atomic<t, addressspace=""> object, T</t,></pre> | operand, memoryOrder). |
| operand, | |
| <pre>memory_order memoryOrder = memory_order::relaxed</pre> | |
|) | |
| <pre>template <typename access::address_space<="" pre="" t,=""></typename></pre> | Deprecated in SYCL 2020. |
| addressSpace> | Equivalent to calling <pre>object.fetch_xor(</pre> |
| <pre>T atomic_fetch_xor(atomic<t, addressspace=""> object, T</t,></pre> | operand, memoryOrder). |
| operand, | |
| <pre>memory_order memoryOrder = memory_order::relaxed</pre> | |
|) | |
| <pre>template <typename access::address_space<="" pre="" t,=""></typename></pre> | Deprecated in SYCL 2020. |
| addressSpace> | Equivalent to calling <pre>object.fetch_min(</pre> |
| <pre>T atomic_fetch_min(atomic<t, addressspace=""> object, T</t,></pre> | operand, memoryOrder). |
| operand, | |
| <pre>memory_order memoryOrder = memory_order::relaxed</pre> | |
|) | |
| <pre>template <typename access::address_space<="" pre="" t,=""></typename></pre> | Deprecated in SYCL 2020. |
| addressSpace> | Equivalent to calling <pre>object.fetch_max(</pre> |
| <pre>T atomic_fetch_max(atomic<t, addressspace=""> object, T</t,></pre> | operand, memoryOrder). |
| operand, | |
| <pre>memory_order memoryOrder = memory_order::relaxed</pre> | |
|) | |
| | End of table |

Table 4.122: Global functions available on atomic types.

4.18 Stream class

The SYCL stream class is a buffered output stream that allows outputting the values of built-in, vector and SYCL types to the console. The implementation of how values are streamed to the console is left as an implementation detail.

The way in which values are output by an instance of the SYCL stream class can also be altered using a range of manipulators.

There are two limits that are relevant for the stream class. The totalBufferSize limit specifies the maximum size of the overall character stream that can be output during a kernel invocation, and the workItemBufferSize limit specifies the maximum size of the character stream that can be output within a work item before a flush must be performed. Both of these limits are specified in bytes. The totalBufferSize limit must be sufficient to contain the characters output by all stream statements during execution of a kernel invocation (the aggregate of outputs from all work items), and the workItemBufferSize limit must be sufficient to contain the characters output within a work item between stream flush operations.

If the totalBufferSize or workItemBufferSize limits are exceeded, it is implementation defined whether the streamed characters exceeding the limit are output, or silently ignored/discarded, and if output it is implementation defined whether those extra characters exceeding the workItemBufferSize limit count toward the totalBufferSize limit. Regardless of this implementation defined behavior of output exceeding the limits, no

undefined or erroneous behavior is permitted of an implementation when the limits are exceeded. Unused characters within workItemBufferSize (any portion of the workItemBufferSize capacity that has not been used at the time of a stream flush) do not count toward the totalBufferSize limit, in that only characters flushed count toward the totalBufferSize limit.

The SYCL stream class provides the common reference semantics (see Section 4.5.3).

4.18.1 Stream class interface

The constructors and member functions of the SYCL stream class are listed in Tables 4.125, 4.126, and 4.127 respectively. The additional common special member functions and common member functions are listed in Tables 4.1 and 4.2, respectively.

The operand types that are supported by the SYCL stream class operator <<() operator are listed in Table 4.123.

The manipulators that are supported by the SYCL stream class operator <<() operator are listed in Table 4.124.

```
1
    namespace sycl {
 2
 3
    enum class stream_manipulator {
 4
      flush,
 5
      dec,
 6
      hex,
 7
      oct,
 8
      noshowbase,
 9
      showbase,
10
      noshowpos,
11
      showpos,
12
      endl,
13
      fixed,
14
      scientific,
15
      hexfloat,
16
      defaultfloat
17
    };
18
19
20
    const stream_manipulator flush = stream_manipulator::flush;
21
22
    const stream_manipulator dec = stream_manipulator::dec;
23
24
    const stream_manipulator hex = stream_manipulator::hex;
25
26
    const stream_manipulator oct = stream_manipulator::oct;
27
28
    const stream_manipulator noshowbase = stream_manipulator::noshowbase;
29
30
    const stream_manipulator showbase = stream_manipulator::showbase;
31
32
    const stream_manipulator noshowpos = stream_manipulator::noshowpos;
33
34
    const stream_manipulator showpos = stream_manipulator::showpos;
35
36
    const stream_manipulator endl = stream_manipulator::endl;
37
```

```
38
   const stream_manipulator fixed = stream_manipulator::fixed;
39
40
    const stream_manipulator scientific = stream_manipulator::scientific;
41
42
    const stream_manipulator hexfloat = stream_manipulator::hexfloat;
43
44
   const stream_manipulator defaultfloat = stream_manipulator::defaultfloat;
45
46
    __precision_manipulator__ setprecision(int precision);
47
    __width_manipulator__ setw(int width);
48
49
50 class stream {
51
    public:
52
53
     stream(size_t totalBufferSize, size_t workItemBufferSize, handler& cgh);
54
      /* -- common interface members -- */
55
56
57
     size_t get_size() const;
58
59
     size_t get_work_item_buffer_size() const;
60
      /* get_max_statement_size() has the same functionality as get_work_item_buffer_size(),
61
         and is provided for backward compatibility. get_max_statement_size() is a deprecated
62
         query. */
63
64
      size_t get_max_statement_size() const;
65 };
66
67
   template <typename T>
68 const stream& operator<<(const stream& os, const T &rhs);</pre>
69
70 } // namespace sycl
```

| Stream operand type | Description |
|--|---|
| char, signed char, unsigned char, int, unsigned int, | Outputs the value as a stream of characters. |
| short, unsigned short, long int, unsigned long int, | |
| long long int, unsigned long long int | |
| float, double, half | Outputs the value according to the precision |
| | of the current statement as a stream of char- |
| | acters. |
| char *, const char * | Outputs the string. |
| T *, const T *, multi_ptr | Outputs the address of the pointer as a |
| | stream of characters. |
| vec | Outputs the value of each component of the |
| | vector as a stream of characters. |
| <pre>id, range, item, nd_item, group, nd_range, h_item</pre> | Outputs the value of each component of |
| | each id or range as a stream of characters. |
| | End of table |

Table 4.123: Operand types supported by the stream class.

| Stream manipulator | Description |
|--------------------|---|
| flush | Triggers a flush operation, which synchro- |
| | nizes the work item stream buffer with the |
| | global stream buffer, and then empties the |
| | work item stream buffer. After a flush, the |
| | full workItemBufferSize is available again |
| | for subsequent streaming within the work |
| | item. |
| endl | Outputs a new-line character and then trig- |
| | gers a flush operation. |
| dec | Outputs any subsequent values in the cur- |
| | rent statement in decimal base. |
| hex | Outputs any subsequent values in the cur- |
| | rent statement in hexadecimal base. |
| oct | Outputs any subsequent values in the cur- |
| | rent statement in octal base. |
| noshowbase | Outputs any subsequent values without the |
| | base prefix. |
| showbase | Outputs any subsequent values with the |
| | base prefix. |
| noshowpos | Outputs any subsequent values without a |
| | plus sign if the value is positive. |
| showpos | Outputs any subsequent values with a plus |
| | sign if the value is positive. |
| setw(int) | Sets the field width of any subsequent val- |
| | ues in the current statement. |
| setprecision(int) | Sets the precision of any subsequent values |
| | in the current statement. |
| fixed | Outputs any subsequent floating-point val- |
| | ues in the current statement in fixed notation. |
| scientific | Outputs any subsequent floating-point val- |
| | ues in the current statement in scientific no- |
| | tation. |
| hexfloat | Outputs any subsequent floating-point val- |
| | ues in the current statement in hexadecimal |
| | notation. |
| defaultfloat | Outputs any subsequent floating-point val- |
| | ues in the current statement in the default no- |
| | tation. |
| | End of table |

Table 4.124: Manipulators supported by the stream class.

| Constructor | Description |
|--|--|
| Constructor | |
| <pre>stream(size_t totalBufferSize, size_t</pre> | Constructs a SYCL stream instance asso- |
| <pre>workItemBufferSize, handler& cgh)</pre> | ciated with the command group specified by |
| | cgh, with a maximum buffer size in bytes per |
| | kernel invocation specified by the parameter |
| | totalBufferSize, and a maximum stream |
| | size that can be buffered by a work item be- |
| | tween stream flushes specified by the param- |
| | eterworkItemBufferSize. |
| | End of table |

Table 4.125: Constructors of the stream class.

| Member function | Description |
|--|--|
| <pre>size_t get_size()const</pre> | Returns the total buffer size, in bytes. |
| <pre>size_t get_work_item_buffer_size()const</pre> | Returns the buffer size per work item, in |
| | bytes. |
| <pre>size_t get_max_statement_size()const</pre> | Deprecated query with same functionality |
| | <pre>as get_work_item_buffer_size().</pre> |
| | End of table |

Table 4.126: Member functions of the stream class.

| Global function | Description |
|---|---|
| <pre>template <typename t=""> const stream& operator<<(const</typename></pre> | Outputs any valid values (see 4.123) as |
| stream& os, const T &rhs) | a stream of characters and applies any valid manipulator (see 4.124) to the current stream. |
| | End of table |

Table 4.127: Global functions of the stream class.

4.18.2 Synchronization

An instance of the SYCL stream class is required to synchronize with the host, and must output everything that is streamed to it via the operator<<() operator before a flush operation (that doesn't exceed the workItemBufferSize or totalBufferSize limits) within a SYCL kernel function by the time that the event associated with a command group submission enters the completed state. The point at which this synchronization occurs and the member function by which this synchronization is performed are implementation defined. For example it is valid for an implementation to use printf().

The SYCL stream class is required to output the content of each stream, between flushes (up to workItemBufferSize), without mixing with content from the same stream in other work items. There are no other output order guarantees between work items or between streams. The stream flush operation therefore delimits the unit of output that is guaranteed to be displayed without mixing with other work items, with respect to a single stream.

4.18.3 Implicit flush

There is guaranteed to be an implicit flush of each stream used by a kernel, at the end of kernel execution, from the perspective of each work item. There is also an implicit flush when the endl stream manipulator is executed. No other implicit flushes are permitted in an implementation.

4.18.4 Performance note

The usage of the stream class is designed for debugging purposes and is therefore not recommended for performance critical applications.

4.19 SYCL built-in functions for SYCL host and device

SYCL kernels may execute on any SYCL device, which requires the functions used in the kernels to be compiled and linked for both device and host. In the SYCL programming model, the built-ins are available for the entire SYCL application within the sycl namespace, although their semantics may be different. This section follows the OpenCL 1.2 specification document [1, ch. 6.12] - except that for SYCL, all functions are located within the sycl namespace - and describes the behavior of these functions for SYCL host and device. The expected precision and any other semantic requirements are defined in the backend specification.

The SYCL built-in functions are available throughout the SYCL application, and depending on where they execute, they are either implemented using their host implementation or the device implementation. The SYCL system guarantees that all of the built-in functions fulfill the same requirements for both host and device.

4.19.1 Description of the built-in types available for SYCL host and device

All of the OpenCL built-in types are available in the namespace sycl. For the purposes of this document we use generic type names for describing sets of valid SYCL types. The generic type names themselves are not valid SYCL types, but they represent a set of valid types, as defined in Tables 4.128. Each generic type within a section is comprised of a combination of scalar, SYCL vec and/or marray class specializations. The letters n and N define valid sizes for class specializations, where n means 2,3,4,8,16 and N means any positive value of size_t type. Note that any reference to the base type refers to the type of a scalar or the element type of a SYCL vec or marray specialization.

| Generic type name | Description |
|-------------------|---|
| floatn | <pre>floatn, mfloatn, marray<n,float></n,float></pre> |
| genfloatf | float, floatn |
| doublen | <pre>doublen, mdoublen, marray<n,double></n,double></pre> |
| genfloatd | double, doublen |
| halfn | <pre>halfn, mhalfn, marray<n,half></n,half></pre> |
| genfloath | half, halfn |
| genfloat | genfloatf, genfloatd, genfloath |
| sgenfloat | float, double, half |
| gengeofloat | <pre>float, float2, float3, float4,</pre> |
| | mfloat2, mfloat3, mfloat4 |
| | Continued on next page |

Table 4.128: Generic type name description, which serves as a description for all valid types of parameters to kernel functions [1].

| Generic type name | Description |
|-------------------|---|
| gengeodouble | <pre>double, double2, double3, double4,</pre> |
| | <pre>mdouble2, mdouble3, mdouble4</pre> |
| charn | <pre>charn, mcharn, marray<n, char=""></n,></pre> |
| scharn | <pre>scharn, mscharn, marray<n,signed char<="" pre=""></n,signed></pre> |
| | > |
| ucharn | ucharn, mucharn, marray <n,unsigned< th=""></n,unsigned<> |
| | char> |
| igenchar | signed char, scharn |
| ugenchar | unsigned char, ucharn |
| genchar | char, charn, igenchar, ugenchar |
| shortn | <pre>shortn, mshortn, marray<n,short></n,short></pre> |
| genshort | short, shortn |
| ushortn | <pre>ushortn, mushortn, marray<n,unsigned< pre=""></n,unsigned<></pre> |
| | short> |
| ugenshort | unsigned short, ushortn |
| uintn | <pre>uintn, muintn, marray<n,unsigned int=""></n,unsigned></pre> |
| ugenint | unsigned int, uintn |
| intn | <pre>intn, mintn, marray<n,int></n,int></pre> |
| genint | int, intn |
| ulongn | <pre>ulongn, mulongn, marray<n,unsigned< pre=""></n,unsigned<></pre> |
| | long int> |
| ugenlong | unsigned long int, ulongn |
| longn | <pre>longn, mlongn, marray<n,long int=""></n,long></pre> |
| genlong | long int, longn |
| ulonglongn | <pre>ulonglongn, mulonglongn, marray<n,< pre=""></n,<></pre> |
| | unsigned long long int> |
| ugenlonglong | unsigned long long int, ulonglongn |
| longlongn | <pre>longlongn, mlonglongn, marray<n,long< pre=""></n,long<></pre> |
| | long int> |
| genlonglong | long long int, longlongn |
| igenlonginteger | genlong, genlonglong |
| ugenlonginteger | ugenlong, ugenlonglong |
| geninteger | genchar, genshort, ugenshort, |
| | genint, ugenint, igenlonginteger, |
| | ugenlonginteger |
| genintegerNbit | All types within geninteger whose |
| | base type are N bits in size, where N |
| | = 8, 16, 32, 64. |
| igeninteger | igenchar, genshort, genint, |
| | igenlonginteger |
| igenintegerNbit | All types within igeninteger whose |
| | base type are N bits in size, where N |
| | = 8, 16, 32, 64. |
| ugeninteger | ugenchar, ugenshort, ugenint, |
| | ugenionginteger |
| | Continued on next page |

Table 4.128: Generic type name description, which serves as a description for all valid types of parameters to kernel functions [1].

| Generic type name | Description |
|-------------------|---|
| ugenintegerNbit | All types within ugeninteger whose |
| | base type are N bits in size, where N |
| | = 8, 16, 32, 64. |
| sgeninteger | char, signed char, unsigned char, |
| | short, unsigned short, int, unsigned |
| | int, long int, unsigned long int, |
| | long long int, unsigned long long int |
| gentype | genfloat, geninteger |
| genfloatptr | All permutations of multi_ptr <datat< th=""></datat<> |
| | , addressSpace, IsDecorated> where |
| | dataT is all types within genfloat, |
| | addressSpace is access::address_space:: |
| | <pre>global_space, access::address_space::</pre> |
| | <pre>local_space and access::address_space</pre> |
| | ::private_space and IsDecorated is |
| | <pre>access::decorated::yes and access::</pre> |
| | decorated::no. |
| genintptr | All permutations of multi_ptr <datat< th=""></datat<> |
| | , addressSpace, IsDecorated> where |
| | dataT is all types within genint, |
| | addressSpace is access::address_space:: |
| | <pre>global_space, access::address_space::</pre> |
| | <pre>local_space and access::address_space</pre> |
| | ::private_space and IsDecorated is |
| | <pre>access::decorated::yes and access::</pre> |
| | decorated::no. |
| booln | <pre>marray<n,bool></n,bool></pre> |
| genbool | bool, booln |
| | End of table |

Table 4.128: Generic type name description, which serves as a description for all valid types of parameters to kernel functions [1].

4.19.2 Work-item functions

In the OpenCL 1.2 specification document [1, ch. 6.12.1] in Table 6.7 the work-item functions are defined where they provide the size of the enqueued kernel NDRange. These functions are available in SYCL through the item and group classes see sections 4.10.1.4, 4.10.1.5 and 4.10.1.7.

4.19.3 Function objects

SYCL provides a number of function objects in the **sycl** namespace on host and device. All function objects obey C++ conversion and promotion rules. Each function object is additionally specialized for **void** as a *transparent* function object that deduces its parameter types and return type.

SYCL function objects can be identified using the sycl::is_native_function_object and sycl::
is_native_function_object_v traits classes.

```
1 namespace sycl {
2
3 template <typename T=void>
4 struct plus {
5
     T operator()(const T& x, const T& y) const;
6 };
7
8 template <typename T=void>
9 struct multiplies {
10
     T operator()(const T& x, const T& y) const;
11 };
12
13 template <typename T=void>
14 struct bit_and {
15
    T operator()(const T& x, const T& y) const;
16 };
17
18 template <typename T=void>
19 struct bit_or {
20
    T operator()(const T& x, const T& y) const;
21 };
22
23 template <typename T=void>
24 struct bit_xor {
25
    T operator()(const T& x, const T& y) const;
26 };
27
28 template <typename T=void>
29
   struct logical_and {
30
    T operator()(const T& x, const T& y) const;
31 };
32
33 template <typename T=void>
34 struct logical_or {
35
    T operator()(const T& x, const T& y) const;
36 };
37
38 template <typename T=void>
39
   struct minimum {
40
    T operator()(const T& x, const T& y) const;
41 };
42
43 template <typename T=void>
44 struct maximum {
45
    T operator()(const T& x, const T& y) const;
46 };
47
48 } // namespace sycl
```

| Member function | Description |
|---|--|
| T operator()(const T& x, const T& y)const | Returns the sum of its arguments, equivalent |
| | to $\mathbf{x} + \mathbf{y}$. |
| | End of table |

Table 4.129: Member functions for the plus function object.

| Member function | Description |
|---|--|
| T operator()(const T& x, const T& y)const | Returns the product of its arguments, equivalent to x * y. |
| | End of table |

Table 4.130: Member functions for the multiplies function object.

| Member function | Description |
|---|---|
| T operator()(const T& x, const T& y)const | Returns the bitwise AND of its arguments, equivalent to x & y. |
| | End of table |

Table 4.131: Member functions for the **bit_and** function object.

| Member function | Description |
|---|--|
| T operator()(const T& x, const T& y)const | Returns the bitwise OR of its arguments, equivalent to $x y$. |
| | End of table |

Table 4.132: Member functions for the **bit_or** function object.

| Member function | Description |
|---|---|
| T operator()(const T& x, const T& y)const | Returns the bitwise XOR of its arguments, equivalent to $\mathbf{x} \cdot \mathbf{y}$. |
| | End of table |

Table 4.133: Member functions for the **bit_xor** function object.

| Member function | Description |
|---|---|
| T operator()(const T& x, const T& y)const | Returns the logical AND of its arguments, equivalent to x && y. |
| | End of table |

Table 4.134: Member functions for the logical_and function object.

| Member function | Description |
|---|--|
| T operator()(const T& x, const T& y)const | Returns the logical OR of its arguments, equivalent to $x \parallel y$. |
| | End of table |

Table 4.135: Member functions for the logical_or function object.

| Member function | Description |
|---|--|
| T operator()(const T& x, const T& y)const | Applies std::less to its arguments, in the same order, then returns the lesser argument unchanged. |
| | End of table |

Table 4.136: Member functions for the minimum function object.

| Member function | Description |
|---|--|
| T operator()(const T& x, const T& y)const | Applies std::greater to its arguments, in the same order, then returns the greater ar- gument unchanged. |
| | End of table |

Table 4.137: Member functions for the maximum function object.

4.19.4 Algorithms library

SYCL provides an algorithms library based on the functions described in Section 28 of the C++17 specification. The first argument to each function is an execution policy, and data ranges are described using instances of multi_ptr (in place of more general iterators) in order to guarantee that address space information is visible to the compiler. The functions defined in this section are free functions available in the sycl namespace.

Any restrictions from the standard algorithms library apply. Some of the functions in the SYCL algorithms library introduce additional restrictions in order to maximize portability across different devices and to minimize the chances of encountering unexpected behavior.

All algorithms are supported for the fundamental scalar types supported by SYCL (see Table 5.1) and instances of the SYCL vec and marray classes. Functions with arguments of type vec<T,N> or marray<T,N> are applied component-wise and are semantically equivalent to N calls to a scalar algorithm with type T.

The execution policy sycl::execution::group denotes that an algorithm should be performed collaboratively by the work-items in the specified group. All algorithms using this execution policy therefore act as group functions (as defined in Section 4.19.5), inheriting all restrictions of group functions. Unless the description of a function says otherwise, how the elements of a range are processed by the work-items in a group is undefined.

4.19.4.1 any_of, all_of and none_of

The any_of, all_of and none_of functions test whether Boolean conditions hold for any of, all of or none of the values in a range, respectively.

| Function | Description |
|---|--|
| <pre>template <typename executionpolicy,="" pre="" ptr,<="" typename=""></typename></pre> | Return true if pred returns true for any |
| typename Predicate> | element in the range [first, last). |
| <pre>bool any_of(ExecutionPolicy&& policy, Ptr first,</pre> | |
| Ptr last, Predicate pred) | If policy is sycl::execution::group, |
| | first and last must be the same for all |
| | work-items in the group; and pred must be |
| | an immutable callable with the same type |
| | and state for all work-items in the group. |
| | End of table |

Table 4.138: Overloads for the any_of function.

| Function | Description |
|---|--|
| <pre>template <typename executionpolicy,="" pre="" ptr,<="" typename=""></typename></pre> | Return true if pred returns true for all |
| typename Predicate> | elements in the range [first, last). |
| <pre>bool all_of(ExecutionPolicy&& policy, Ptr first,</pre> | |
| Ptr last, Predicate pred) | If policy is sycl::execution::group, |
| | first and last must be the same for all |
| | work-items in the group; and pred must be |
| | an immutable callable with the same type |
| | and state for all work-items in the group. |
| | End of table |

Table 4.139: Overloads for the all_of function.

| Function | Description |
|---|--|
| <pre>template <typename executionpolicy,="" pre="" ptr,<="" typename=""></typename></pre> | Return true if pred returns true for no |
| typename Predicate> | elements in the range [first, last). |
| <pre>bool none_of(ExecutionPolicy&& policy, Ptr first</pre> | |
| , Ptr last, Predicate pred) | If policy is sycl::execution::group, |
| | first and last must be the same for all |
| | work-items in the group; and pred must be |
| | an immutable callable with the same type |
| | and state for all work-items in the group. |
| | End of table |

Table 4.140: Overloads for the none_of function.

4.19.4.2 reduce

The reduce function combines values in an unspecified order using a binary operator. The result of a call to reduce is non-deterministic if the binary operator is not commutative and associative. Only the binary operators defined in Section 4.19.3 are supported by reduce in SYCL 2020, but the standard C++ syntax is used for forward compatibility with future SYCL versions.

| Function | Description |
|---|--|
| <pre>template <typename executionpolicy,="" pre="" ptr,<="" typename=""></typename></pre> | Combine the values in the range [<i>first</i> , <i>last</i>) |
| typename BinaryOperation> | using the operator binary_op, which must |
| <pre>Ptr::value_type reduce(ExecutionPolicy&& policy,</pre> | be an instance of a SYCL function object. |
| Ptr first, Ptr last, BinaryOperation binary_op) | <pre>binary_op(*first, *first) must return a</pre> |
| | value of type Ptr::value_type. |
| | |
| | If policy is sycl::execution::group, |
| | first, last and the type of binary_op must |
| | be the same for all work-items in the group. |
| <pre>template <typename executionpolicy,="" pre="" ptr,<="" typename=""></typename></pre> | Combine the values in the range [<i>first</i> , <i>last</i>) |
| typename T, typename BinaryOperation> | using an initial value of init and the opera- |
| <pre>T reduce(ExecutionPolicy&& policy, Ptr first,</pre> | tor binary_op, which must be an instance of |
| Ptr last, T init, BinaryOperation binary_op) | a SYCL function object. binary_op(init, |
| | *first) must return a value of type T. |
| | |
| | If policy is <pre>sycl::execution::group,</pre> |
| | first, last, init and the type of binary_op |
| | must be the same for all work-items in the |
| | group. |
| | End of table |

Table 4.141: Overloads of the reduce function.

4.19.4.3 exclusive_scan and inclusive_scan

The scan functions compute a generalized prefix sum using a binary operator. The result of a call to a scan is non-deterministic if the binary operator is not associative. Only the binary operators defined in Section 4.19.3 are supported by the scan functions in SYCL 2020, but the standard C++ syntax is used for forward compatibility with future SYCL versions.

A scan operation can be exclusive or inclusive. For a scan of elements $[x_0, ..., x_n]$, the *i*th result in an exclusive scan excludes x_i , whereas the *i*th result in an inclusive scan includes x_i .

| Function | Description |
|---|--|
| <pre>template <typename executionpolicy,="" inptr,<="" pre="" typename=""></typename></pre> | Perform an exclusive scan over the values |
| typename OutPtr, typename BinaryOperation> | in the range [<i>first</i> , <i>last</i>) using the operator |
| <pre>OutPtr exclusive_scan(ExecutionPolicy&& policy,</pre> | binary_op, which must be an instance of |
| InPtr first, InPtr last, OutPtr result, | a SYCL function object. binary_op(* |
| BinaryOperation binary_op) | <pre>first, *first) must return a value of type OutPtr::value_type.</pre> |
| | The value written to result $+ i$ is the exclusive scan of the first <i>i</i> values in the range and the identity value of binary_op. Returns a pointer to the end of the output range. |
| | If policy is sycl::execution::group, first, last, result and the type of binary_op must be the same for all work-items in the group. |
| <pre>template <typename executionpolicy,="" inptr,<="" pre="" typename=""></typename></pre> | Perform an exclusive scan over the values |
| typename OutPtr, typename T, typename | in the range [first, last) using the operator |
| BinaryOperation> | binary_op, which must be an instance of a |
| <pre>OutPtr exclusive_scan(ExecutionPolicy&& policy,</pre> | SYCL function object. binary_op(init, * |
| InPtr first, InPtr last, OutPtr result, T init, BinaryOperation binary_op) | first) must return a value of type T. |
| | The value written to result $+ i$ is the exclusive scan of the first i values in the range and an initial value specified by init. Returns a pointer to the end of the output range. |
| | If policy is sycl::execution::group, |
| | first, last, result, init and the type |
| | of binary_op must be the same for all |
| | work-items in the group. |
| | End of table |

Table 4.142: Overloads of the exclusive_scan function.

| Function | Description |
|---|---|
| <pre>template <typename executionpolicy,="" inptr,<="" pre="" typename=""></typename></pre> | Perform an inclusive scan over the values |
| <pre>typename OutPtr, typename BinaryOperation></pre> | in the range [first, last) using the operator |
| <pre>OutPtr inclusive_scan(ExecutionPolicy&& policy,</pre> | binary_op, which must be an instance of |
| InPtr first, InPtr last, OutPtr result, | a SYCL function object. binary_op(* |
| BinaryOperation binary_op) | <pre>first, *first) must return a value of type OutPtr::value_type.</pre> |
| | The value written to result $+ i$ is the inclusive scan of the first <i>i</i> values in the range. Returns a pointer to the end of the output range. |
| | If policy is sycl::execution::group, first, last, result and the type of binary_op must be the same for all work-items in the group. |
| <pre>template <typename executionpolicy,="" inptr,<="" pre="" typename=""></typename></pre> | Perform an inclusive scan over the values |
| <pre>typename OutPtr, typename BinaryOperation, typename</pre> | in the range [first, last) using the operator |
| T> | binary_op, which must be an instance of a |
| <pre>OutPtr inclusive_scan(ExecutionPolicy&& policy,</pre> | SYCL function object. binary_op(init, * |
| InPtr first, InPtr last, OutPtr result, | first) must return a value of type T. |
| BinaryOperation binary_op, T init) | The value written to result $+ i$ is the inclusive scan of the first <i>i</i> values in the range and an initial value specified by init. Returns a pointer to the end of the output range. |
| | If policy is such consention comment |
| | first last result init and the type |
| | of hinary on must be the same for all |
| | work-items in the group. |
| | End of table |

Table 4.143: Overloads of the inclusive_scan function.

4.19.5 Group functions

SYCL provides a number of functions that expose functionality tied to groups of work-items (such as group barriers and collective operations). These group functions act as synchronization points and must be encountered in converged control flow by all work-items in the group — if one work-item in the group reaches the function, then all work-items in the group must reach the function. Additionally, restrictions may be placed on the arguments passed to each function in order to ensure that all work-items in the group agree on the operation that is being performed. Any such restrictions on the arguments passed to a function are defined within the descriptions of those functions. Violating these restrictions results in undefined behavior.

The functions defined in this section are free functions available in the sycl namespace, and they all accept a templated parameter of type Group. It is valid to pass an argument of type group or of type sub_group for this parameter.

All group functions are supported for the fundamental scalar types supported by SYCL (see Table 5.1) and instances of the SYCL vec and marray classes. Functions with arguments of type vec<T,N> or marray<T,N> are applied component-wise and are semantically equivalent to N calls to a scalar function with type T.

Using a group function inside of a kernel may introduce additional limits on the resources available to user code inside the same kernel. The behavior of these limits is implementation-defined, but must be reflected by calls to kernel querying functions (such as kernel::get_info) as described in Section 4.12.

It is undefined behavior for any group function to be invoked within a parallel_for_work_group or parallel_for_work_item context.

4.19.5.1 group_broadcast

The group_broadcast function communicates a value held by one work-item to all other work-items in the group.

| Function | Description |
|--|--|
| template <typename group,="" t="" typename=""></typename> | Broadcast the value of x from the work-item |
| T group_broadcast(Group g, T x) | with the smallest linear id to all work-items |
| | within the group. |
| <pre>template <typename group,="" t="" typename=""></typename></pre> | Broadcast the value of x from the work-item |
| <pre>T group_broadcast(Group g, T x, Group::</pre> | with the specified linear id to all work-items |
| linear_id_type local_linear_id) | within the group. |
| | |
| | The value of local_linear_id must be |
| | the same for all work-items in the group. |
| template <typename group,="" t="" typename=""></typename> | Broadcast the value of x from the work-item |
| T group_broadcast(Group g, T x, Group::id_type | with the specified id to all work-items |
| local_id) | within the group. |
| | |
| | The value of local_id must be the |
| | same for all work-items in the group, |
| | and its dimensionality must match the |
| | dimensionality of the group. |
| | End of table |

Table 4.144: Overloads of the group_broadcast function.

4.19.5.2 group_barrier

The group_barrier function synchronizes all work-items in a group, using a group barrier.

| Function | Description |
|---|--|
| <pre>template <typename group=""></typename></pre> | Synchronizes all work-items in the group. |
| <pre>void group_barrier(Group g, memory_scope</pre> | The current work-item will wait at the |
| <pre>fence_scope = Group::fence_scope)</pre> | barrier until all work-items in the group |
| | have reached the barrier. In addition, |
| | the barrier performs a group mem-fence |
| | operation ensuring that all memory accesses |
| | issued before the barrier complete before |
| | those issued after the barrier: all work-items |
| | in the group execute a release fence prior |
| | to synchronizing at the barrier, and all |
| | work-items in the group execute an acquire |
| | fence afterwards. |
| | |
| | By default, the scope of these fences is |
| | set to the narrowest scope including all |
| | work-items in the group (as reported by |
| | Group::fence_scope). This scope may be |
| | optionally overridden with a broader scope, |
| | specified by the fence_scope argument. |
| | End of table |

Table 4.145: Overloads for the group_barrier function.

4.19.5.3 group_any_of, group_all_of and group_none_of

The group_any_of, group_all_of and group_none_of functions correspond to the any_of, all_of and none_of functions from the algorithm library in Section 4.19.4, respectively. The group_ variants of the functions perform the same operations, but apply directly to values supplied by the work-items in a group instead of a range of values stored in memory.

| Function | Description |
|--|---|
| <pre>template <typename group=""></typename></pre> | Return true if pred is true for any work-item |
| <pre>bool group_any_of(Group g, bool pred)</pre> | in the group. |
| template <typename group,="" t,="" td="" typename="" typename<=""><td>Return true if pred(x) is true for any</td></typename> | Return true if pred(x) is true for any |
| Predicate> | work-item in the group. |
| <pre>bool group_any_of(Group g, T x, Predicate pred)</pre> | |
| | pred must be an immutable callable with |
| | the same type and state for all work-items |
| | in the group. |
| | End of table |

Table 4.146: Overloads for the group_any_of function.

| Function | Description |
|--|--|
| <pre>template <typename group=""></typename></pre> | Return true if pred is true for all work-items |
| <pre>bool group_all_of(Group g, bool pred)</pre> | in the group. |
| | Continued on next page |

Table 4.147: Overloads for the group_all_of function.

| Function | Description |
|--|--|
| template <typename group,="" t,="" td="" typename="" typename<=""><td>Return true if pred(x) is true for all work-</td></typename> | Return true if pred(x) is true for all work- |
| Predicate> | items in the group. |
| <pre>bool group_all_of(Group g, T x, Predicate pred)</pre> | |
| | pred must be an immutable callable with |
| | the same type and state for all work-items |
| | in the group. |
| | End of table |

Table 4.147: Overloads for the group_all_of function.

| Function | Description |
|---|--|
| <pre>template <typename group=""></typename></pre> | Return true if pred is true for no work-item |
| <pre>bool group_none_of(Group g, bool pred)</pre> | in the group. |
| template <typename group,="" t,="" td="" typename="" typename<=""><td>Return true if pred(x) is true for no work-</td></typename> | Return true if pred(x) is true for no work- |
| Predicate> | item in the group. |
| <pre>bool group_none_of(Group g, T x, Predicate pred)</pre> | |
| | pred must be an immutable callable with |
| | the same type and state for all work-items |
| | in the group. |
| | End of table |

Table 4.148: Overloads for the group_none_of function.

4.19.5.4 group_reduce

The group_reduce function corresponds to the reduce function from the algorithms library in Section 4.19.4. The group_ variant of the function performs the same operation, but applies directly to values supplied by the work-items in a group instead of a range of values stored in memory.

| Function | Description |
|---|---|
| template <typename group,="" t,="" td="" typename="" typename<=""><td>Combine the values of x from all work-items</td></typename> | Combine the values of x from all work-items |
| BinaryOperation> | in the group using the operator binary_op, |
| T group_reduce(Group g, T x, BinaryOperation | which must be an instance of a SYCL |
| binary_op) | function object. binary_op(x, x) must |
| | return a value of type T. |
| | |
| | The type of binary_op must be the |
| | same for all work-items in the group. |
| template <typename group,="" t,<="" td="" typename="" v,=""><td>Combine the values of x from all work-items</td></typename> | Combine the values of x from all work-items |
| typename BinaryOperation> | in the group using an initial value of init |
| T group_reduce(Group g, V x, T init, | and the operator binary_op, which must |
| BinaryOperation binary_op) | be an instance of a SYCL function object. |
| | binary_op(init, x) must return a value of |
| | type T. |
| | |
| | The type of binary_op must be the |
| | same for all work-items in the group. |
| | End of table |
4.19.5.5 group_exclusive_scan and group_inclusive_scan

The group_exclusive_scan and group_inclusive_scan functions correspond to the exclusive_scan and inclusive_scan functions from the algorithm library in Section 4.19.4, respectively. The group_ variants of the functions perform the same operations, but apply directly to values supplied by the work-items in a group instead of a range of values stored in memory.

| Function | Description |
|---|--|
| template <typename group,="" t,="" td="" typename="" typename<=""><td>Perform an exclusive scan over the values</td></typename> | Perform an exclusive scan over the values |
| BinaryOperation> | of x from all work-items in the group using |
| <pre>T group_exclusive_scan(Group g, T x,</pre> | the operator binary_op, which must be |
| BinaryOperation binary_op) | an instance of a SYCL function object. |
| | binary_op(x, x) must return a value of type T. |
| | The value returned on work-item i is the exclusive scan of the first i work-items in the group and the identity value of binary_op. For multi-dimensional groups, the order of work-items in the group is determined by their linear id. |
| | The type of binary_op must be the same for all work-items in the group. |
| <pre>template <typename group,="" pre="" t,<="" typename="" v,=""></typename></pre> | Perform an exclusive scan over the values |
| typename BinaryOperation> | of x from all work-items in the group using |
| T group_exclusive_scan(Group g, V x, T init, | the operator binary_op, which must be |
| BinaryOperation binary_op) | an instance of a SYCL function object. |
| | type T. |
| | The value returned on work-item i is the exclusive scan of the first i work items in the group and an initial value specified by init. For multi-dimensional groups, the order of work-items in the group is determined by their linear id. |
| | init and the type of binary_op must |
| | be the same for all work-items in the group. |
| | End of table |

Table 4.150: Overloads of the group_exclusive_scan function.

| Function | Description |
|---|--|
| template <typename group,="" t,="" td="" typename="" typename<=""><td>Perform an inclusive scan over the values</td></typename> | Perform an inclusive scan over the values |
| BinaryOperation> | of x from all work-items in the group using |
| T group_inclusive_scan(Group g, T x, | the operator binary_op, which must be |
| BinaryOperation binary_op) | an instance of a SYCL function object. |
| | binary_op(x, x) must return a value of type T. |
| | The value returned on work-item i is |
| | the inclusive scan of the first <i>i</i> work items |
| | in the group. For multi-dimensional groups, |
| | the order of work-items in the group is |
| | determined by their linear ld. |
| | The type of binary_op must be the same for all work-items in the group. |
| template <typename group,="" td="" typename="" typename<="" v,=""><td>Perform an inclusive scan over the values</td></typename> | Perform an inclusive scan over the values |
| BinaryOperation, typename T> | of \mathbf{x} from all work-items in the group using |
| T group_inclusive_scan(Group g, V x, | the operator binary_op, which must be |
| BinaryOperation binary_op, T init) | an instance of a SYCL function object. |
| | binary_op(init, x) must return a value of |
| | type 1. |
| | The value returned on work-item i is the inclusive scan of the first i work items |
| | in the group and an initial value specified |
| | by init. For multi-dimensional groups. |
| | the order of work-items in the group is |
| | determined by their linear id. |
| | |
| | init and the type of binary_op must |
| | be the same for all work-items in the group. |
| | End of table |

Table 4.151: Overloads of the group_inclusive_scan function.

4.19.6 Math functions

In SYCL the OpenCL math functions are available in the namespace sycl on host and device with the same precision guarantees as defined in the OpenCL 1.2 specification document [1, ch. 7] for host and device. For a SYCL platform the numerical requirements for host need to match the numerical requirements of the OpenCL math built-in functions. The built-in functions can take as input float or optionally double and their vec and marray counterparts, for all supported dimensions including dimension 1.

The built-in functions available for SYCL host and device, with the same precision requirements for both host and device, are described in Table 4.152.

| Math Function | Description |
|--|---|
| genfloat acos (genfloat x) | Inverse cosine function. |
| genfloat acosh (genfloat x) | Inverse hyperbolic cosine. |
| genfloat acospi (genfloat x) | Compute $\arccos(x)/\pi$ |
| genfloat asin (genfloat x) | Inverse sine function. |
| genfloat asinh (genfloat x) | Inverse hyperbolic sine. |
| genfloat asinpi (genfloat x) | Compute $\arcsin(x)/\pi$ |
| <pre>genfloat atan (genfloat y_over_x)</pre> | Inverse tangent function. |
| genfloat atan2 (genfloat y, genfloat x) | Compute $\arctan(y/x)$. |
| genfloat atanh (genfloat x) | Hyperbolic inverse tangent. |
| genfloat atanpi (genfloat x) | Compute $\arctan(x)/\pi$. |
| genfloat atan2pi (genfloat y, genfloat x) | Compute $\arctan 2(y, x)/\pi$. |
| <pre>genfloat cbrt (genfloat x)</pre> | Compute cube-root. |
| <pre>genfloat ceil (genfloat x)</pre> | Round to integral value using the round to |
| | positive infinity rounding mode. |
| <pre>genfloat copysign (genfloat x, genfloat y)</pre> | Returns x with its sign changed to match the |
| | sign of y. |
| <pre>genfloat cos (genfloat x)</pre> | Compute cosine. |
| <pre>genfloat cosh (genfloat x)</pre> | Compute hyperbolic cosine. |
| genfloat cospi (genfloat x) | Compute $\cos(\pi x)$. |
| <pre>genfloat erfc (genfloat x)</pre> | Complementary error function. |
| <pre>genfloat erf (genfloat x)</pre> | Error function encountered in integrating the |
| | normal distribution. |
| <pre>genfloat exp (genfloat x)</pre> | Compute the base- e exponential of x. |
| <pre>genfloat exp2 (genfloat x)</pre> | Exponential base 2 function. |
| <pre>genfloat exp10 (genfloat x)</pre> | Exponential base 10 function. |
| <pre>genfloat expm1 (genfloat x)</pre> | Compute $e^x - 1.0$. |
| genfloat fabs (genfloat x) | Compute absolute value of a floating-point |
| | number. |
| genfloat fdim (genfloat x, genfloat y) | x - y if $x > y, +0$ if x is less than or equal to |
| | у. |
| genfloat floor (genfloat x) | Round to integral value using the round to |
| | negative infinity rounding mode. |
| genfloat fma (genfloat a, genfloat b, genfloat c) | Returns the correctly rounded floating-point |
| | representation of the sum of c with the in- |
| | inducts and b. Round- |
| | ang of intermediate products shall not oc- |
| | cur. Edge case behavior is per the IEEE /34- |
| confloat fmax (confloat x confloat x) | 2008 statuard. Returns y if $x < y$ otherwise it returns y |
| genfloat fmax (genfloat x, genfloat y) | If one argument is a NaN fmax() returns |
| geniioat imax (geniioat x, sgeniioat y) | the other argument. If both arguments are |
| | NaNs fmax() returns a NaN |
| genfloat fmin (genfloat x genfloat x) | Returns v if $v < x$ otherwise it returns x |
| genfloat fmin (genfloat \mathbf{x} , genfloat \mathbf{y}) | If one argument is a NaN $fmin()$ returns |
| generout init (generout it, sychirout j) | the other argument. If both arguments are |
| | NaNs, fmin() returns a NaN. |
| genfloat fmod (genfloat x, genfloat v) | Modulus. Returns $x - v \cdot trunc(x/v)$. |
| , <u>journa</u> , j | Continued on next page |

Table 4.152: Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL1.2 specification [1].CHAPTER 4. SYCL PROGRAMMING INTERFACE363

| Math Function | Description |
|---|--|
| genfloat fract (genfloat x, genfloatptr iptr) | Returns fmin(x – floor (x), |
| | nextafter(genfloat(1.0), genfloat(0.0))). |
| | floor(x) is returned in iptr. |
| <pre>genfloat frexp (genfloat x, genintptr exp)</pre> | Extract mantissa and exponent from x. For |
| | each component the mantissa returned is a |
| | float with magnitude in the interval $[1/2, 1)$ |
| | or 0. Each component of x equals mantissa |
| | returned $\times 2^{exp}$. |
| genfloat hypot (genfloat x, genfloat y) | Compute the value of the square root of x^2 + |
| | y2 without undue overflow or underflow. |
| genint ilogb (genfloat x) | Return the exponent as an integer value. |
| genfloat Idexp (genfloat x, genint k) | Multiply x by 2 to the power k. |
| genfloat Idexp (genfloat x, int k) | La contra Data a dia a dia |
| genfloat Igamma (genfloat x) | Log gamma function. Returns the natu- |
| | rai logarithin of the absolute value of the |
| genfloat laamma r (genfloat y genintetr signa) | I og gamma function. Returns the natu |
| geniioat igamma_i (geniioat x, genincpti signp) | ral logarithm of the absolute value of the |
| | gamma function. The sign of the gamma |
| | function is returned in the sign of the gamma |
| | lgamma r. |
| genfloat log (genfloat x) | Compute natural logarithm. |
| genfloat log2 (genfloat x) | Compute a base 2 logarithm. |
| genfloat log10 (genfloat x) | Compute a base 10 logarithm. |
| <pre>genfloat log1p (genfloat x)</pre> | Compute $\log_e(1.0 + x)$. |
| genfloat logb (genfloat x) | Compute the exponent of x, which is the |
| | integral part of logr (x) . |
| genfloat mad (genfloat a,genfloat b, genfloat c) | mad approximates a $*$ b + c. Whether or |
| | how the product of a * b is rounded and |
| | how supernormal or subnormal intermediate |
| | products are handled is not defined. mad is |
| | intended to be used where speed is preferred |
| confloat marmag (confloat a confloat a) | Over accuracy. |
| geniioat maxmag (geniioat x, geniioat y) | $\begin{array}{c} \text{Keturns x in } x \ge y , \text{ y in } y \ge x , \text{ otherwise} \\ \text{fmax}(x, y) \end{array}$ |
| genfloat minmag (genfloat y genfloat y) | $\frac{1}{2} = \frac{1}{2} $ |
| genitoat minimag (genitoat x, genitoat y) | fmin(x y) |
| genfloat modf (genfloat x, genfloatptr intr) | Decompose a floating-point number. The |
| generout moul (generout ii, generoutput iput) | modf function breaks the argument x into in- |
| | tegral and fractional parts, each of which has |
| | the same sign as the argument. It stores the |
| | integral part in the object pointed to by iptr. |
| genfloatf nan (ugenint nancode) | Returns a quiet NaN. The nancode may be |
| genfloatd nan (ugenlonginteger nancode) | placed in the significand of the resulting |
| | NaN. |
| | Continued on next page |

Table 4.152: Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1].

| Math Function | Description |
|--|--|
| <pre>genfloat nextafter (genfloat x, genfloat y)</pre> | Computes the next representable single- |
| | precision floating-point value following x in |
| | the direction of y. Thus, if y is less than x, |
| | nextafter() returns the largest representable |
| | floating-point number less than x. |
| genfloat pow (genfloat x, genfloat y) | Compute x to the power y. |
| genfloat pown (genfloat x, genint y) | Compute x to the power y, where y is an |
| | integer. |
| genfloat powr (genfloat x, genfloat y) | Compute x to the power y, where $x \ge 0$. |
| genfloat remainder (genfloat x, genfloat y) | Compute the value r such that $r = x - n^*y$, |
| | where n is the integer nearest the exact value |
| | of x/y. If there are two integers closest to x/y, |
| | n shall be the even one. If r is zero, it is given |
| | The same sign as x. |
| genfloat remquo (genfloat x, genfloat y, genintptr | The remultion function computes the value r such that $n = n$, $h^* x_1$, where h is the inter- |
| quo) | such that $f = x - k^{+}y$, where k is the inte- |
| | get nearest the exact value of x/y . If there are two integers closest to x/y k shall be the |
| | even one. If r is zero, it is given the same |
| | sign as x. This is the same value that is re- |
| | turned by the remainder function remain |
| | also calculates the lower seven bits of the in- |
| | tegral quotient x/y , and gives that value the |
| | same sign as x/v . It stores this signed value |
| | in the object pointed to by quo. |
| genfloat rint (genfloat x) | Round to integral value (using round to |
| | nearest even rounding mode) in floating- |
| | point format. Refer to section 7.1 of the |
| | OpenCL 1.2 specification document [1] for |
| | description of rounding modes. |
| genfloat rootn (genfloat x, genint y) | Compute x to the power 1/y. |
| <pre>genfloat round (genfloat x)</pre> | Return the integral value nearest to x round- |
| | ing halfway cases away from zero, regard- |
| | less of the current rounding direction. |
| genfloat rsqrt (genfloat x) | Compute inverse square root. |
| genfloat sin (genfloat x) | Compute sine. |
| genfloat sincos (genfloat x, genfloatptr cosval) | Compute sine and cosine of x. The com- |
| | puted sine is the return value and computed |
| renfleet einh (renfleet n) | Compute hyperbolic sine |
| genfloat sinni (genfloat x) | Compute hyperbolic sine. |
| genfloat shipi (genfloat x) | Compute sill (π, χ) . |
| genfloat tan (genfloat x) | Compute tangent |
| genfloat tank (genfloat x) | Compute hyperbolic tangent |
| genfloat tanni (genfloat x) | Compute hyperbolic tangent. |
| genfloat toamma (genfloat x) | Compute the gamma function |
| gentrout (gentrout x) | Continued on next page |

Table 4.152: Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1].

CHAPTER 4. SYCL PROGRAMMING INTERFACE

| Math Function | Description |
|-----------------------------|--|
| genfloat trunc (genfloat x) | Round to integral value using the round to |
| | zero rounding mode. |
| | End of table |

Table 4.152: Math functions which work on SYCL host and device. They correspond to Table 6.8 of the OpenCL 1.2 specification [1].

In SYCL the implementation defined precision math functions are defined in the namespace sycl::native. The functions that are available within this namespace are specified in Tables 4.153.

| Native Math Function | Description |
|--|--|
| <pre>genfloatf cos (genfloatf x)</pre> | Compute cosine over an implementation- |
| | defined range. The maximum error is |
| | implementation-defined. |
| <pre>genfloatf divide (genfloatf x, genfloatf y)</pre> | Compute x / y over an implementation- |
| | defined range. The maximum error is |
| | implementation-defined. |
| <pre>genfloatf exp (genfloatf x)</pre> | Compute the base- e exponential of x over |
| | an implementation-defined range. The max- |
| | imum error is implementation-defined. |
| <pre>genfloatf exp2 (genfloatf x)</pre> | Compute the base- 2 exponential of x over |
| | an implementation-defined range. The max- |
| | imum error is implementation-defined. |
| <pre>genfloatf exp10 (genfloatf x)</pre> | Compute the base- 10 exponential of x over |
| | an implementation-defined range. The max- |
| | imum error is implementation-defined. |
| <pre>genfloatf log (genfloatf x)</pre> | Compute natural logarithm over an imple- |
| | mentation defined range. The maximum er- |
| | ror is implementation-defined. |
| <pre>genfloatf log2 (genfloatf x)</pre> | Compute a base 2 logarithm over an |
| | implementation-defined range. The maxi- |
| | mum error is implementation-defined. |
| <pre>genfloatf log10 (genfloatf x)</pre> | Compute a base 10 logarithm over an |
| | implementation-defined range. The maxi- |
| | mum error is implementation-defined. |
| <pre>genfloatf powr (genfloatf x, genfloatf y)</pre> | Compute x to the power y, where |
| | $x \ge 0$. The range of x and y are |
| | implementation-defined. The maximum er- |
| | ror is implementation-defined. |
| genfloatf recip (genfloatf x) | Compute reciprocal over an |
| | implementation-defined range. The |
| | maximum error is implementation-defined. |
| geniloati rsqrt (geniloati x) | Compute inverse square root over an |
| | implementation-defined range. The maxi- |
| | mum error is implementation-defined. |
| | Continued on next page |

Table 4.153: Native math functions.

CHAPTER 4. SYCL PROGRAMMING INTERFACE

| Native Math Function | Description |
|---|--|
| <pre>genfloatf sin (genfloatf x)</pre> | Compute sine over an implementation- |
| | defined range. The maximum error is |
| | implementation-defined. |
| <pre>genfloatf sqrt (genfloatf x)</pre> | Compute square root over an |
| | implementation-defined range. The |
| | maximum error is implementation-defined. |
| <pre>genfloatf tan (genfloatf x)</pre> | Compute tangent over an implementation- |
| | defined range. The maximum error is |
| | implementation-defined. |
| | End of table |

Table 4.153: Native math functions.

In SYCL the half precision math functions are defined in sycl::half_precision. The functions that are available within this namespace are specified in Tables 4.154. These functions are implemented with a minimum of 10-bits of accuracy i.e. an ULP value is less than or equal to 8192 ulp.

| Half Math function | Description |
|--|---|
| <pre>genfloatf cos (genfloatf x)</pre> | Compute cosine. x must be in the range - |
| | 216 to +216. |
| <pre>genfloatf divide (genfloatf x, genfloatf y)</pre> | Compute x / y. |
| <pre>genfloatf exp (genfloatf x)</pre> | Compute the base- e exponential of x. |
| <pre>genfloatf exp2 (genfloatf x)</pre> | Compute the base- 2 exponential of x. |
| <pre>genfloatf exp10 (genfloatf x)</pre> | Compute the base- 10 exponential of x. |
| <pre>genfloatf log (genfloatf x)</pre> | Compute natural logarithm. |
| <pre>genfloatf log2 (genfloatf x)</pre> | Compute a base 2 logarithm. |
| <pre>genfloatf log10 (genfloatf x)</pre> | Compute a base 10 logarithm. |
| <pre>genfloatf powr (genfloatf x, genfloatf y)</pre> | Compute x to the power y, where $x \ge 0$. |
| <pre>genfloatf recip (genfloatf x)</pre> | Compute reciprocal. |
| <pre>genfloatf rsqrt (genfloatf x)</pre> | Compute inverse square root. |
| <pre>genfloatf sin (genfloatf x)</pre> | Compute sine. x must be in the range -216 |
| | to +216. |
| <pre>genfloatf sqrt (genfloatf x)</pre> | Compute square root. |
| <pre>genfloatf tan (genfloatf x)</pre> | Compute tangent. x must be in the range - |
| | 216 to +216. |
| | End of table |

4.19.7 Integer functions

Integer math functions are available in SYCL in the namespace sycl on host and device. The built-in functions can take as input char, unsigned char, short, unsigned short, int, unsigned int, long long int, unsigned long long int and their vec and marray counterparts. The supported integer math functions are described in Table 4.155.

| Integer Function | Description |
|--|--|
| ugeninteger abs (geninteger x) | Returns $ x $. |
| ugeninteger abs_diff (geninteger x, geninteger y) | Returns $ x - y $ without modulo overflow. |
| <pre>geninteger add_sat (geninteger x, geninteger y)</pre> | Returns $x + y$ and saturates the result. |
| geninteger hadd (geninteger x, geninteger y) | Returns $(x + y) >> 1$. The intermediate sum |
| | does not modulo overflow. |
| geninteger rhadd (geninteger x, geninteger y) | Returns $(x + y + 1) >> 1$. The intermediate |
| | sum does not modulo overflow. |
| geninteger clamp (geninteger x, geninteger minval, | Returns min(max(x, minval), maxval). Re- |
| geninteger maxval) | sults are undefined if minval > maxval. |
| geninteger clamp (geninteger x, sgeninteger | |
| minval, sgeninteger maxval) | |
| geninteger clz (geninteger x) | Returns the number of leading 0-bits in x, |
| | starting at the most significant bit position. |
| | If x is 0, returns the size in bits of the type |
| | of x or component type of x, if x is a vector |
| | lype. |
| geninteger ctz (geninteger x) | x is 0, returns the size in hits of the type of y |
| | or component type of x if x is a vector type |
| conjuteger mad hi (| Beturns mult bi $(a, b) + c$ |
| geninteger a geninteger b geninteger c) | Keturns mur_nr(a, b)+c. |
| geninteger mad sat (geninteger a | Returns a $*$ b + c and saturates the result |
| geninteger h. geninteger c) | Roturns a s + c and saturates the result. |
| geninteger max (geninteger x, geninteger y) | Returns v if $x < v$, otherwise it returns x. |
| geninteger max (geninteger x, sgeninteger y) | 5 57 |
| geninteger min (geninteger x, geninteger y) | Returns y if $y < x$, otherwise it returns x. |
| geninteger min (geninteger x, sgeninteger y) | |
| geninteger mul_hi (geninteger x, geninteger y) | Computes x * y and returns the high half |
| | of the product of x and y. |
| geninteger rotate (geninteger v, geninteger i) | For each element in v, the bits are shifted |
| | left by the number of bits given by the corre- |
| | sponding element in 1 (subject to usual shift |
| | modulo rules described in section 6.3). Bits |
| | shifted off the left side of the element are |
| | Shifted back in from the right. |
| geninteger sub_sat (geninteger x, geninteger y) | Returns $x - y$ and saturates the result. |
| ugenintegeribbit upsample (ugeninteger8bit ni, | result[1] = ((usnort)n1[1] << 8) 10[|
| igeninteger16hit unsample (igeninteger8hit hi | $\frac{1}{1}$ |
| ugenintegeribbit lo) | |
| ugeninteger32bit upsample (ugeninteger16bit hi. | result[i] = ((uint)hi[i] << 16) lo[i |
| ugeninteger16bit lo) |] |
| igeninteger32bit upsample (igeninteger16bit hi | result[i] = ((int)hi[i] << 16) lo[i] |
| ugeninteger16bit lo) | |
| ugeninteger64bit upsample (ugeninteger32bit hi, | <pre>result[i] = ((ulonglong)hi[i] << 32) </pre> |
| ugeninteger32bit lo) | lo[i] |
| | Continued on next page |

Table 4.155: Integer functions which work on SYCL host and device, are available in the sycl namespace.

CHAPTER 4. SYCL PROGRAMMING INTERFACE

| Integer Function | Description |
|--|---|
| igeninteger64bit upsample (igeninteger32bit hi, | <pre>result[i] = ((longlong)hi[i] << 32) </pre> |
| ugeninteger32bit lo) | lo[i] |
| geninteger popcount (geninteger x) | Returns the number of non-zero bits in x. |
| geninteger32bit mad24 (geninteger32bit x, | Multiply two 24-bit integer values x and y |
| geninteger32bit y, geninteger32bit z) | and add the 32-bit integer result to the 32- |
| | bit integer z. Refer to definition of mul24 to |
| | see how the 24-bit integer multiplication is |
| | performed. |
| <pre>geninteger32bit mul24 (geninteger32bit x,</pre> | Multiply two 24-bit integer values x and y. |
| geninteger32bit y) | x and y are 32-bit integers but only the low |
| | 24-bits are used to perform the multiplica- |
| | tion. mul24 should only be used when val- |
| | ues in x and y are in the range $[-2^{23}, 2^{23}-1]$ |
| | if x and y are signed integers and in the range |
| | $[0, 2^{24} - 1]$ if x and y are unsigned integers. |
| | If x and y are not in this range, the multipli- |
| | cation result is implementation-defined. |
| | End of table |

Table 4.155: Integer functions which work on SYCL host and device, are available in the sycl namespace.

4.19.8 **Common functions**

In SYCL the OpenCL common functions are available in the namespace sycl on host and device as defined in the OpenCL 1.2 specification document [1, par. 6.12.4]. They are described here in Table 4.156. The built-in functions can take as input float or optionally double and their vec and marray counterparts.

| Common Function | Description |
|--|--|
| <pre>genfloat clamp (genfloat x, genfloat minval,</pre> | Returns fmin(fmax(x, minval), maxval). Re- |
| genfloat maxval) | sults are undefined if <i>minval</i> > <i>maxval</i> . |
| <pre>genfloatf clamp (genfloatf x, float minval, float</pre> | |
| maxval) | |
| <pre>genfloatd clamp (genfloatd x, double minval,</pre> | |
| double maxval) | |
| genfloat degrees (genfloat radians) | Converts radians to degrees, i.e. $(180/\pi) *$ |
| | radians. |
| genfloat max (genfloat x, genfloat y) | Returns y if $x < y$, otherwise it returns x. If |
| <pre>genfloatf max (genfloatf x, float y)</pre> | x or y are infinite or NaN, the return values |
| <pre>genfloatd max (genfloatd x, double y)</pre> | are undefined. |
| genfloat min (genfloat x, genfloat y) | Returns y if $y < x$, otherwise it returns x. If |
| <pre>genfloatf min (genfloatf x, float y)</pre> | x or y are infinite or NaN, the return values |
| <pre>genfloatd min (genfloatd x, double y)</pre> | are undefined. |
| genfloat mix (genfloat x, genfloat y, genfloat a) | Returns the linear blend of x&y imple- |
| <pre>genfloatf mix (genfloatf x, genfloatf y, float a)</pre> | mented as: $x+(y-x)*a$. <i>a</i> must be a value in |
| genfloatd mix (genfloatd x, genfloatd y, double | the range 0.0 1.0. If a is not in the range |
| a) | 0.0 1.0, the return values are undefined. |
| | Continued on next page |

Table 4.156: Common functions which work on SYCL host and device, are available in the sycl namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1]. CHAPTER 4. SYCL PROGRAMMING INTERFACE 369

| Common Function | Description |
|---|--|
| genfloat radians (genfloat degrees) | Converts degrees to radians, i.e. $(\pi/180) *$ |
| | degrees. |
| <pre>genfloat step (genfloat edge, genfloat x)</pre> | Returns 0.0 if $x < edge$, otherwise it returns |
| <pre>genfloatf step (float edge, genfloatf x)</pre> | 1.0. |
| <pre>genfloatd step (double edge, genfloatd x)</pre> | |
| <pre>genfloat smoothstep (genfloat edge0, genfloat edge1,</pre> | Returns 0.0 if $x \le edge0$ and 1.0 if $x \ge edge0$ |
| genfloat x) | edge1 and performs smooth Hermite inter- |
| <pre>genfloatf smoothstep (float edge0, float edge1,</pre> | polation between 0 and 1 when $edge0 <$ |
| genfloatf x) | x < edge1. This is useful in cases where |
| <pre>genfloatd smoothstep (double edge0, double edge1</pre> | you would want a threshold function with a |
| , genfloatd x) | smooth transition. |
| | This is equivalent to: |
| | gentype t; |
| | t = clamp ((x <= edge 0)/ (edge1 >= |
| | edge0), 0, 1); |
| | return t * t * (3 - 2 * t); |
| | |
| | Results are undefined if $edge0 >= edge1$ or |
| | if x, edge0 or edge1 is a NaN. |
| genfloat sign (genfloat x) | Returns 1.0 if $x > 0$, -0.0 if $x = -0.0$, $+0.0$ |
| | if $x = +0.0$, or -1.0 if $x < 0$. Returns 0.0 if |
| | x is a NaN. |
| | End of table |

Table 4.156: Common functions which work on SYCL host and device, are available in the sycl namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1].

4.19.9 Geometric functions

In SYCL the OpenCL *geometric* functions are available in the namespace sycl on host and device as defined in the OpenCL 1.2 specification document [1, par. 6.12.5]. The built-in functions can take as input float or optionally double and their vec and codeinlinemarray counterparts, for dimensions 2, 3 and 4. On the host the vector types use the vec class and on an SYCL device use the corresponding native SYCL backend vector types. All of the geometric functions use round-to-nearest-even rounding mode. Table 4.157 contains the definitions of supported geometric functions.

| Geometric Function | Description |
|--|---|
| float4 cross (float4 p0, float4 p1) | Returns the cross product of p0.xyz and |
| float3 cross (float3 p0, float3 p1) | p1.xyz. The <i>w</i> component of float4 result |
| double4 cross (double4 p0, double4 p1) | returned will be 0.0. |
| double3 cross (double3 p0, double3 p1) | |
| mfloat4 cross (mfloat4 p0, mfloat4 p1) | Returns the cross product of first 3 compo- |
| mfloat3 cross (mfloat3 p0, mfloat3 p1) | nents of p0 and p1. The 4th component of |
| mdouble4 cross (mdouble4 p0, mdouble4 p1) | result returned will be 0.0. |
| mdouble3 cross (mdouble3 p0, mdouble3 p1) | |
| <pre>float dot (gengeofloat p0, gengeofloat p1)</pre> | Compute dot product. |
| <pre>double dot (gengeodouble p0, gengeodouble p1)</pre> | |
| | Continued on next page |

Table 4.157: Geometric functions which work on SYCL host and device, are available in the sycl namespace. They correspond to Table 6.13 of the OpenCL 1.2 specification of the Reference of the sycl namespace of the system of th

| Geometric Function | Description |
|---|--|
| <pre>float distance (gengeofloat p0, gengeofloat p1)</pre> | Returns the distance between p0 and p1. |
| <pre>double distance (gengeodouble p0, gengeodouble p1)</pre> | This is calculated as $length(p0 - p1)$. |
| <pre>float length (gengeofloat p)</pre> | Return the length of vector p, i.e., |
| double length (gengeodouble p) | $\sqrt{p.x^2 + p.y^2 + \dots}$ |
| <pre>gengeofloat normalize (gengeofloat p)</pre> | Returns a vector in the same direction as p |
| gengeodouble normalize (gengeodouble p) | but with a length of 1. |
| <pre>float fast_distance (gengeofloat p0, gengeofloat p1)</pre> | Returns fast_length(p0 - p1). |
| <pre>float fast_length (gengeofloat p)</pre> | Returns the length of vector p computed |
| | <pre>as: sqrt((half)(pow(p.x,2)+ pow(p.y,2)</pre> |
| | +)) |
| | <pre>computed as: p * rsqrt((half)(pow(p.x ,2)+ pow(p.y,2)+)) The result shall be within 8192 ulps error from the infinitely precise result of if (all (p == 0.0f)) result = p; else result = p / sqrt (pow(p.x,2)+ pow(p. y,2)+); with the following exceptions: 1. If the sum of squares is greater than FLT_MAX then the value of the floating- point values in the result vector are</pre> |
| | a. If the sum of squares is less than FLT_MIN then the implementation may return back p. 3. If the device is in "denorms are flushed to zero" mode, individual operand elements with magnitude less than sqrt(FLT_MIN) may be flushed to zero before proceeding with the calculation. |
| | End of table |

Table 4.157: Geometric functions which work on SYCL host and device, are available in the sycl namespace. They correspond to Table 6.13 of the OpenCL 1.2 specification [1].

4.19.10 Relational functions

In SYCL the OpenCL *relational* functions are available in the namespace sycl on host and device as defined in the OpenCL 1.2 specification document [1, par. 6.12.6]. The built-in functions can take as input char, unsigned char, short, unsigned short, int, unsigned int, long, unsigned long, float or optionally double and their vec and marray counterparts. The relational functions are provided in addition to the the operators.

The available built-in functions for vec template class are described in Tables 4.158

CHAPTER 4. SYCL PROGRAMMING INTERFACE

| Relational Function | Description |
|--|--|
| <pre>igeninteger32bit isequal (genfloatf x, genfloatf y)</pre> | Returns the component-wise compare of |
| <pre>igeninteger64bit isequal (genfloatd x, genfloatd y)</pre> | x == y. |
| <pre>igeninteger32bit isnotequal (genfloatf x, genfloatf</pre> | Returns the component-wise compare of |
| y) | x! = y. |
| <pre>igeninteger64bit isnotequal (genfloatd x, genfloatd</pre> | |
| y) | |
| igeninteger32bit isgreater (genfloatf x, genfloatf y | Returns the component-wise compare of $x >$ |
|) | у. |
| igeninteger64bit isgreater (genfloatd x, genfloatd y | |
| | Defense (hereinen er einen e |
| igeninteger32bit isgreaterequal (genfloatf x, | Returns the component-wise compare of |
| geniloati y) | $x \ge y$. |
| igeninteger64bit isgreaterequal (genfloatd x, | |
| igeninteger22bit isloss (genfloatf y genfloatf y) | Returns the component wise compare of $r < r$ |
| igeninteger64bit isless (genfloatd x, genfloatd y) | x < x < y |
| igeninteger32hit islessequal (genfloatf x genfloatf | y. Returns the component-wise compare of |
| v) | $r \le v$ |
| igeninteger64bit islessequal (genfloatd x, genfloatd | $x \leq y$. |
| v) | |
| igeninteger32bit islessgreater (genfloatf x, | Returns the component-wise compare of |
| genfloatf y) | (x < y) (x > y). |
| <pre>igeninteger64bit islessgreater (genfloatd x,</pre> | |
| genfloatd y) | |
| <pre>igeninteger32bit isfinite (genfloatf x)</pre> | Test for finite value. |
| igeninteger64bit isfinite (genfloatd x) | |
| <pre>igeninteger32bit isinf (genfloatf x)</pre> | Test for infinity value (positive or negative). |
| igeninteger64bit isinf (genfloatd x) | |
| <pre>igeninteger32bit isnan (genfloatf x)</pre> | Test for a NaN. |
| igeninteger64bit isnan (genfloatd x) | |
| igeninteger32bit isnormal (genfloatf x) | lest for a normal value. |
| igeninteger64bit isnormal (genfloatd X) | Test if any ments are and and isondered() |
| igeninteger32bit isordered (geniloati X, geniloati Y | takes arguments x and x and returns the re |
| J | sult is equal $(x, x) \& \& w$ is equal (y, y) |
|) | suit isequal(x, x) && isequal(y, y). |
| igeninteger32bit isunordered (genfloatf x, genfloatf | Test if arguments are unordered. |
| v) | isunordered() takes arguments x and y. |
| igeninteger64bit isunordered (genfloatd x. genfloatd | returning non-zero if x or y is NaN, and |
| y) | zero otherwise. |
| <pre>igeninteger32bit signbit (genfloatf x)</pre> | Test for sign bit. The scalar version of the |
| <pre>igeninteger64bit signbit (genfloatd x)</pre> | function returns a 1 if the sign bit in the float |
| | is set else returns 0. |
| | The vector version of the function returns |
| | the following for each component in <i>floatn</i> : |
| | -1 (i.e all bits set) if the sign bit in the float |
| | is set else returns 0. |
| | Continued on next page |

Table 4.158: Relational functions for vec template class which work on SYCL host and device, are available in the sycl namespace. They correspond to Table 6.14 of the OpenCL 1.2 specification [1]. 372

CHAPTER 4. SYCL PROGRAMMING INTERFACE

| Relational Function | Description |
|--|--|
| int any (igeninteger x) | Returns 1 if the most significant bit in any |
| | component of x is set; otherwise returns 0. |
| int all (igeninteger x) | Returns 1 if the most significant bit in all |
| | components of x is set; otherwise returns 0. |
| gentype bitselect (gentype a, gentype b, gentype c) | Each bit of the result is the corresponding |
| | bit of a if the corresponding bit of c is 0. |
| | Otherwise it is the corresponding bit of b. |
| geninteger select (geninteger a, geninteger b, | For each component of a vector type: |
| igeninteger c) | <pre>result[i] = (MSB of c[i] is set)? b[i</pre> |
| geninteger select (geninteger a, geninteger b, |] : a[i]. |
| ugeninteger c) | For a scalar type: |
| <pre>genfloatf select (genfloatf a, genfloatf b,</pre> | result = c ? b : a. |
| genint c) | geninteger must have the same number of |
| <pre>genfloatf select (genfloatf a, genfloatf b,</pre> | elements and bits as gentype. |
| ugenint c) | |
| genfloatd select (genfloatd a, genfloatd b, | |
| igeninteger64 c) | |
| genfloatd select (genfloatd a, genfloatd b, | |
| ugeninteger64 c) | |
| | End of table |

Table 4.158: Relational functions for vec template class which work on SYCL host and device, are available in the sycl namespace. They correspond to Table 6.14 of the OpenCL 1.2 specification [1].

| Relational Function | Description |
|--|---|
| <pre>genbool isequal (genfloatf x, genfloatf y)</pre> | Returns the component-wise compare of |
| <pre>genbool isequal (genfloatd x, genfloatd y)</pre> | x == y. |
| <pre>genbool isnotequal (genfloatf x, genfloatf y)</pre> | Returns the component-wise compare of |
| <pre>genbool isnotequal (genfloatd x, genfloatd y)</pre> | x! = y. |
| <pre>genbool isgreater (genfloatf x, genfloatf y)</pre> | Returns the component-wise compare of $x >$ |
| <pre>genbool isgreater (genfloatd x, genfloatd y)</pre> | у. |
| <pre>genbool isgreaterequal (genfloatf x, genfloatf y)</pre> | Returns the component-wise compare of |
| <pre>genbool isgreaterequal (genfloatd x, genfloatd y)</pre> | $x \ge y.$ |
| <pre>genbool isless (genfloatf x, genfloatf y)</pre> | Returns the component-wise compare of $x <$ |
| <pre>genbool isless (genfloatd x, genfloatd y)</pre> | у. |
| <pre>genbool islessequal (genfloatf x, genfloatf y)</pre> | Returns the component-wise compare of |
| <pre>genbool islessequal (genfloatd x, genfloatd y)</pre> | $x \ll y$. |
| <pre>genbool islessgreater (genfloatf x, genfloatf y)</pre> | Returns the component-wise compare of |
| <pre>genbool islessgreater (genfloatd x, genfloatd y)</pre> | (x < y) (x > y). |
| <pre>genbool isfinite (genfloatf x)</pre> | Test for finite value. |
| <pre>genbool isfinite (genfloatd x)</pre> | |
| <pre>genbool isinf (genfloatf x)</pre> | Test for infinity value (positive or negative). |
| <pre>genbool isinf (genfloatd x)</pre> | |
| genbool isnan (genfloatf x) | Test for a NaN. |
| genbool isnan (genfloatd x) | |
| | Continued on next page |

The available built-in functions for marray template class are described in Tables 4.159

 Table 4.159:
 Relational functions for scalar data types and marray template class template class which work on

 EXAPPER and Strice preservation the NTE properties.
 373

| Relational Function | Description |
|---|---|
| <pre>genbool isnormal (genfloatf x)</pre> | Test for a normal value. |
| <pre>genbool isnormal (genfloatd x)</pre> | |
| <pre>genbool isordered (genfloatf x, genfloatf y)</pre> | Test if arguments are ordered. isordered() |
| <pre>genbool isordered (genfloatd x, genfloatd y)</pre> | takes arguments x and y, and returns the re- |
| | sult isequal(x, x) && isequal(y, y). |
| <pre>genbool isunordered (genfloatf x, genfloatf y)</pre> | Test if arguments are unordered. |
| <pre>genbool isunordered (genfloatd x, genfloatd y)</pre> | isunordered() takes arguments x and y, |
| | returning true if x or y is NaN, and false |
| | otherwise. |
| <pre>genbool signbit (genfloatf x)</pre> | Test for sign bit, returning true if the sign |
| <pre>genbool signbit (genfloatd x)</pre> | bit in the float is set, and false otherwise. |
| bool any (genbool x) | Returns true if the most significant bit in |
| | any component of x is set; otherwise returns |
| | false. |
| int all (igeninteger x) | Returns true if the most significant bit in |
| | all components of x is set; otherwise returns |
| | false. |
| gentype bitselect (gentype a, gentype b, gentype c) | Each bit of the result is the corresponding |
| | bit of a if the corresponding bit of c is 0. |
| | Otherwise it is the corresponding bit of b. |
| <pre>gentype select (gentype a, gentype b, genbool c)</pre> | Returns the component-wise result = c ? |
| | b : a. |
| | End of table |

Table 4.159: Relational functions for scalar data types and marray template class template class which work on SYCL host and device, are available in the sycl namespace.

4.19.11 Vector data load and store functions

The functionality from the OpenCL functions as defined in the OpenCL 1.2 specification document [1, par. 6.12.7] is available in SYCL through the vec class in section 4.16.2.

4.19.12 Synchronization functions

In SYCL the OpenCL *synchronization* functions are available through the nd_item class 4.10.1.5, as they are applied to work-items for local or global address spaces. Please see 4.83.

4.19.13 printf function

The functionality of the printf function is covered by the stream class 4.18, which has the capability to print to standard output all of the SYCL classes and primitives, and covers the capabilities defined in the OpenCL 1.2 specification document [1, par. 6.12.13].

5. SYCL Device Compiler

This section specifies the requirements of the SYCL device compiler. Most features described in this section relate to underlying SYCL backend capabilities of target devices and limiting the requirements of device code to ensure portability.

5.1 Offline compilation of SYCL source files

There are two alternatives for a SYCL device compiler: a *single-source device compiler* and a device compiler that supports the technique of SMCP.

A SYCL device compiler takes in a C++ source file, extracts only the SYCL kernels and outputs the device code in a form that can be enqueued from host code by the associated SYCL runtime. How the SYCL runtime invokes the kernels is implementation defined, but a typical approach is for a device compiler to produce a header file with the compiled kernel contained within it. By providing a command-line option to the host compiler, it would cause the implementation's SYCL header files to #include the generated header file. The SYCL specification has been written to allow this as an implementation approach in order to allow SMCP. However, any of the mechanisms needed from the SYCL compiler, the SYCL runtime and build system are implementation defined, as they can vary depending on the platform and approach.

A SYCL single-source device compiler takes in a C++ source file and compiles both host and device code at the same time. This specification specifies how a SYCL single-source device compiler sees and outputs device code for kernels, but does not specify the host compilation.

5.2 Naming of kernels

SYCL kernels are extracted from C++ source files and stored in an implementation- defined format. In the case of the shared-source compilation model, the kernels have to be uniquely identified by both host and device compiler. This is required in order for the host runtime to be able to load the kernel by using the OpenCL host runtime interface.

From this requirement the following rules apply for naming the kernels:

- The kernel name is a *C*++ *typename*.
- The kernel name type may not be forward declared other than in namespace scope (including global namespace scope). If it isn't forward declared but is specified as a template argument in a kernel invoking interface, as described in 4.10.7, then it may not conflict with a name in any enclosing namespace scope.
- If the kernel is defined as a named function object type, the name can be the typename of the function object as long as it is either declared at namespace scope, or does not conflict with any name in an enclosing namespace scope.
- If the kernel is defined as a lambda, a typename can optionally be provided to the kernel invoking interface as described in 4.10.7, so that the developer can control the kernel name for purposes such as debugging or

referring to the kernel when applying build options.

In both single-source and shared-source implementations, a device compiler should detect the kernel invocations (e.g. parallel_for<kernelname>) in the source code and compile the enclosed kernels, storing them with their associated type name.

The format of the kernel and the compilation techniques are implementation defined. The interface between the compiler and the runtime for extracting and executing SYCL kernels on the device is implementation defined.

5.3 Compilation of functions

The SYCL device compiler parses an entire C++ source file supplied by the user, including any header files referenced via **#include** directives. From this source file, the SYCL device compiler must compile kernels for the device, as well as any functions that the kernels call.

The device compiler identifies kernels by looking for calls to kernel invocation commands such as parallel_for. One of the parameters is a function object which is known as a SYCL kernel function, and this function must always return void. Any function called by the SYCL kernel function is also compiled for the device, and these functions together with the SYCL kernel functions are known as device functions. The device compiler searches recursively for any functions called from a device function, and these functions are also compiled for the device and known as device functions.

To illustrate, the following source code shows three functions and a kernel invoke with comments explaining which functions need to be compiled for the device.

```
1
   void f(handler& cgh) {
 2
      // Function "f" is not compiled for device
3
4
      cgh.single_task([=] {
        // This code is compiled for device
5
        g(); // This line forces "g" to be compiled for device
6
 7
      });
8 }
9
10 void g() {
      // Called from kernel, so "g" is compiled for device
11
12 }
13
14 void h() {
15
     // Not called from a device function, so not compiled for device
16 }
```

In order for the SYCL device compiler to correctly compile device functions, all functions in the source file, whether device functions or not, must be syntactically correct functions according to this specification. A syntactically correct function adheres to at least the minimum required C++ version defined in Section 3.8.1.

5.4 Language restrictions for device functions

Device functions must abide by certain restrictions. The full set of C++ features are not available to these functions. Following is a list of these restrictions:

CHAPTER 5. SYCL DEVICE COMPILER

376

- Structures containing pointers may be shared. However, when a pointer is passed between SYCL devices or between the host and a SYCL device, dereferencing that pointer on the device produces undefined behavior unless the device supports USM and the pointer is an address within a USM memory region (see Section 4.8).
- Memory storage allocation is not allowed in kernels. All memory allocation for the device is done on the host using accessor classes or using USM as explained in Section 4.8. Consequently, the default allocation operator new overloads that allocate storage are disallowed in a SYCL kernel. The placement new operator and any user-defined overloads that do not allocate storage are permitted.
- Kernel functions must always have a void return type. A kernel lambda trailing-return-type that is not void is therefore illegal, as is a return statement (that would return from the kernel function) with an expression that does not convert to void.
- The odr-use of polymorphic classes and classes with virtual inheritance is allowed. However, no virtual member functions are allowed to be called in a SYCL kernel or any functions called by the kernel.
- No function pointers or references are allowed to be called in a SYCL kernel or any functions called by the kernel.
- RTTI is disabled inside kernels.
- No variadic functions are allowed to be called in a SYCL kernel or any functions called by the kernel.
- Exception-handling cannot be used inside a SYCL kernel or any code called from the kernel. But of course noexcept is allowed.
- Recursion is not allowed in a SYCL kernel or any code called from the kernel.
- Variables with thread storage duration (thread_local storage class specifier) are not allowed to be odr-used in kernel code.
- Variables with static storage duration that are odr-used inside a kernel must be const or constexpr and zero-initialized or constant-initialized.
- The rules for kernels apply to both the kernel function objects themselves and all functions, operators, member functions, constructors and destructors called by the kernel. This means that kernels can only use library functions that have been adapted to work with SYCL. Implementations are not required to support any library routines in kernels beyond those explicitly mentioned as usable in kernels in this spec. Developers should refer to the SYCL built-in functions in 4.19 to find functions that are specified to be usable in kernels.
- Interacting with a special SYCL runtime class (i.e. SYCL accessor, sampler or stream) that is stored within a C++ union is undefined behavior.

5.5 Built-in scalar data types

In a SYCL device compiler, the device definition of all standard C++ fundamental types from Table 5.1 must match the host definition of those types, in both size and alignment. A device compiler may have this preconfigured so that it can match them based on the definitions of those types on the platform, or there may be a necessity for a device compiler command-line option to ensure the types are the same.

CHAPTER 5. SYCL DEVICE COMPILER

| Fundamental data type | Description |
|------------------------|---|
| bool | A conditional data type which can be ei- |
| | ther true or false. The value true expands |
| | to the integer constant 1 and the value false |
| | expands to the integer constant 0. |
| char | A signed or unsigned 8-bit integer, as de- |
| | fined by the C++ core language |
| signed char | A signed 8-bit integer, as defined by the |
| | C++ core language |
| unsigned char | An unsigned 8-bit integer, as defined by the |
| | C++ core language |
| short int | A signed integer of at least 16-bits, as de- |
| | fined by the C++ core language |
| unsigned short int | An unsigned integer of at least 16-bits, as |
| | defined by the C++ core language |
| int | A signed integer of at least 16-bits, as de- |
| | fined by the C++ core language |
| unsigned int | An unsigned integer of at least 16-bits, as |
| | defined by the C++ core language |
| long int | A signed integer of at least 32-bits, as de- |
| | fined by the C++ core language |
| unsigned long int | An unsigned integer of at least 32-bits, as |
| | defined by the C++ core language |
| long long int | An integer of at least 64-bits, as defined by |
| | the C++ core language |
| unsigned long long int | An unsigned integer of at least 64-bits, as |
| | defined by the C++ core language |
| float | A 32-bit floating-point. The float data type |
| | must conform to the IEEE 754 single preci- |
| | sion storage format. |
| double | A 64-bit floating-point. The double data |
| | type must conform to the IEEE 754 double |
| | precision storage format. |
| | End of table |

The standard C++ fixed width types, e.g. int8_t, int16_t, int32_t, int64_t, should have the same size as defined by the C++ standard for host and device.

Table 5.1: Fundamental data types supported by SYCL.

5.6 Preprocessor directives and macros

The standard C++ preprocessing directives and macros are supported. The following preprocessor macros must be defined by all conformant implementations:

• SYCL_LANGUAGE_VERSION substitutes an integer reflecting the version number and revision of the SYCL language being supported by the implementation. The version of SYCL defined in this document will have SYCL_LANGUAGE_VERSION substitute the integer 202001, composed with the general SYCL version followed by 2 digits representing the revision number;

- __SYCL_DEVICE_ONLY__ is defined to 1 if the source file is being compiled with a SYCL device compiler which does not produce host binary;
- __SYCL_SINGLE_SOURCE__ is defined to 1 if the source file is being compiled with a SYCL single-source compiler which produces host as well as device binary;
- SYCL_EXTERNAL is an optional macro which enables external linkage of SYCL functions and member functions to be included in a SYCL kernel. The macro is only defined if the implementation supports external linkage. For more details see 5.9.1

In addition, for each SYCL backend supported, the preprocessor macros described in the backend section 4.1 must be defined by all conformant implementations.

5.7 Kernel attributes

The SYCL general programming interface defines attributes that augment the information available while generating the device code for a particular platform.

5.7.1 Core kernel attributes

The attributes in Table 5.2 are defined in the [[sycl::]] namespace and are applied to the function-type of kernel function declarations using C++ attribute specifier syntax.

A given attribute-token shall appear at most once in each attribute-list. The first declaration of a function shall specify an attribute if any declaration of that function specifies the same attribute. If a function is declared with an attribute in one translation unit and the same function is declared without the same attribute in another translation unit, the program is ill-formed and no diagnostic is required.

If there are any conflicts between different kernel attributes, then the behavior is undefined. The attributes have an effect when applied to a kernel function and no effect otherwise (i.e. no effect on non-kernel functions and on anything other than a function). If an attribute is applied to a device function that is not a kernel function (but that is potentially called from a kernel function), then the effect is implementation defined. It is implementation defined whether any diagnostic is produced when an attribute is applied to anything other than the function-type of a kernel function.

| SYCL attribute | Description |
|---|---|
| <pre>reqd_work_group_size(dim0)</pre> | Indicates that the kernel must be launched |
| <pre>reqd_work_group_size(dim0, dim1)</pre> | with the specified work-group size. The |
| <pre>reqd_work_group_size(dim0, dim1, dim2)</pre> | order of the arguments matches the con- |
| | structor of the group class. Each argument |
| | to the attribute must be an integral con- |
| | stant expression. The dimensionality of |
| | the attribute variant used must match the |
| | dimensionality of the work-group used to |
| | invoke the kernel. |
| | SYCL device compilers should give a |
| | compilation error if the required work- |
| | group size is unsupported. If the kernel |
| | is submitted for execution using an in- |
| | compatible work-group size, the SYCL |
| | runtime must throw an exception with the |
| | errc::nd_range_error error code. |
| <pre>work_group_size_hint(dim0)</pre> | Hint to the compiler on the work-group size |
| <pre>work_group_size_hint(dim0, dim1)</pre> | most likely to be used when launching the |
| <pre>work_group_size_hint(dim0, dim1, dim2)</pre> | kernel at runtime. Each argument must be an |
| | integral constant expression, and the number |
| | of dimensional values defined provide ad- |
| | ditional information to the compiler on the |
| | dimensionality most likely to be used when |
| | launching the kernel at runtime. The effect |
| | of this attribute, if any, is implementation |
| | defined. |
| vec_type_hint(<type>)</type> | Hint to the compiler on the vector computa- |
| | tional width of of the kernel. The argument |
| | must be one of the vector types defined in |
| | section 4.10.2. The effect of this autibute, if |
| | any, is implementation defined. |
| | This attribute is deprecated (available |
| | for use, but will likely be removed in a |
| | future version of the specification and is not |
| | recommended for use in new code). |
| | Continued on next page |

Table 5.2: Attributes supported by the SYCL General programming interface.

| SYCL attribute | Description |
|--------------------------|--|
| reqd_sub_group_size(dim) | Indicates that the kernel must be compiled |
| | and executed with the specified sub-group |
| | size. The argument to the attribute must be |
| | an integral constant expression. |
| | |
| | SYCL device compilers should give a |
| | compilation error if the required sub-group |
| | size is unsupported by the device or incom- |
| | patible with any language feature used by |
| | the kernel. The set of valid sub-group sizes |
| | for a kernel can be queried as described in |
| | Table 4.19 and Table 4.99. |
| | End of table |

Table 5.2: Attributes supported by the SYCL General programming interface.

Other attributes may be provided as part of SYCL backend-interop functionality.

5.7.2 Example attribute syntax

Using [[sycl::reqd_work_group_size(16)]] as an example attribute, but applying equally to all attributes in Table 5.2 and to attributes that are part of SYCL backend-interop or extensions, the following code examples demonstrate how to apply the attributes to the function-type of a kernel function.

```
1 // Kernel defined as a lambda
2
    myQueue.submit([&](handler &h) {
3
    h.parallel_for( range<1>(16),
4
          [=] (item<1> it) [[sycl::reqd_work_group_size(16)]] {
            //[kernel code]
5
6
          });
7
    });
8
9
   // Kernel defined as a functor to be invoked later
10 class KernelFunctor {
11
      public:
12
      void operator()(item<1> it) const [[sycl::reqd_work_group_size(16)]] {
        //[kernel code]
13
14
      };
15 };
```

5.7.3 Deprecated attribute syntax

The SYCL 1.2.1 specification (superseded by this version) defined two mechanisms for kernel attributes to be specified, which are deprecated in this version of SYCL. The old syntaxes are supported but will be removed in a future version, and are therefore not recommended for use. Specifically, the following two attribute syntaxes defined by the SYCL 1.2.1 specification are deprecated:

1. The attribute syntax defined by the OpenCL C specification within device code (__attribute__((attrib))).

CHAPTER 5. SYCL DEVICE COMPILER

2. The C++ attribute specifier syntax in the [[cl::]] namespace applied to device functions (not the function-type of a kernel function), including automatic propagation of the attribute to any caller of such device functions.

5.8 Address-space deduction

C++ has no type-level support to represent address spaces. As a consequence, the SYCL generic programming model does not directly affect the C++ type of unannotated pointers and references.

Source level guarantees about addresse spaces in the SYCL generic programming model can only be achieved using pointer classes (instances of multi_ptr), which are regular classes that represent pointers to data stored in the corresponding address spaces.

In SYCL, the address space of pointer and references are derived from:

- Accessors that give access to shared data. They can be bound to a memory object in a command group and passed into a kernel. Accessors are used in scheduling of kernels to define ordering. Accessors to buffers have a compile-time address space based on their access mode.
- Explicit pointer classes (e.g. global_ptr) holds a pointer which is known to be addressing the address space represented by the access::address_space. This allows the compiler to determine whether the pointer references global, local, constant or private memory and generate code accordingly.
- Raw C++ pointer and reference types (e.g. int*) are allowed within SYCL kernels. They can be constructed from the address of local variables, explicit pointer classes, or accessors.

5.8.1 Address space assignment

In order to understand where data lives, the device compiler is expected to assign address spaces while lowering types for the underlying target based on the context. The address space deducing rules differ slightly depending on the target device of the SYCL backend.

If the target of the SYCL backend can represent the generic address space, then the "common address space deduction rules" (section 5.8.2) and the "generic as default address space rules" (section 5.8.3) apply. If the target of the SYCL backend cannot represent the generic address space, then the "common address space deduction rules" (section 5.8.2) and the "inferred address space rules" (section 5.8.4) apply.

[Note: SYCL address space does not affect the type, address space shall be understood as memory segment in which data is allocated. For instance, if int i; is allocated to the global address space, then decltype(&i) shall evaluate to int*. — end note]

5.8.2 Common address space deduction rules

The variable declarations get assigned to an address space depending on their scope and storage class:

- Namespace scope
 - The declaration is assigned to global address space if the type is not const
 - The declaration is assigned to constant address space if the type is const
- Block scope and function parameter scope

- Declarations with static storage are treated the same way as variables in namespace scope
- Otherwise the declaration is assigned to the local address space if declared in a hierarchical context
- Otherwise the declaration is assigned to the private address space
- Class scope:
 - Static data members are treated the same way as for variable in namespace scope

The result of a prvalue-to-xvalue conversion is assigned to the local address space if it happens in a hierarchical context or to the private address space otherwise.

5.8.3 Generic as default address space

Unannotated pointers and references are considered to be pointing to the generic address space.

5.8.4 Inferred address space

[Note for this provisional version: *The address space deduction feature described next is being reworked to better align with addition of generic address space and generic as default address space.* — end note]

Inside kernels, the SYCL device compiler will need to auto-deduce the memory region of unannotated pointer and reference types during the lowering of types from C++ to the underlying representation.

If a kernel function or device function contains a pointer or reference type, then the address space deduction must be attempted using the following rules:

- If an explicit pointer class is converted into a C++ pointer value, then the C++ pointer value will point to same address space as the one represented by the explicit pointer class.
- If a variable is declared as a pointer type, but initialized in its declaration to a pointer value with an alreadydeduced address space, then that variable will have the same address space as its initializer.
- If a function parameter is declared as a pointer type, and the argument is a pointer value with a deduced address space, then the function will be compiled as if the parameter had the same address space as its argument. It is legal for a function to be called in different places with different address spaces for its arguments: in this case the function is said to be "duplicated" and compiled multiple times. Each duplicated instance of the function must compile legally in order to have defined behavior.
- If a function return type is declared as a pointer type and return statements use address space deduced expressions, then the function will be compiled as if the return type had the same address space. To compile legally, all return expressions must deduce to the same address space.
- The rules for pointer types also apply to reference types. i.e. a reference variable takes its address space from its initializer. A function with a reference parameter takes its address space from its argument.
- If no other rule above can be applied to a declaration of a pointer, then it is assumed to be in the private address space.

It is illegal to assign a pointer value addressing one address space to a pointer variable addressing a different address space.

CHAPTER 5. SYCL DEVICE COMPILER

5.9 SYCL offline linking

5.9.1 SYCL functions and member functions linkage

The default behavior in SYCL applications is that all the definitions and declarations of the functions and member functions are available to the SYCL compiler, in the same translation unit. When this is not the case, all the symbols that need to be exported to a SYCL library or from a C++ library to a SYCL application need to be defined using the macro: SYCL_EXTERNAL.

The SYCL_EXTERNAL macro will only be defined if the implementation supports offline linking. The macro is implementation-defined, but the following restrictions apply:

- SYCL_EXTERNAL can only be used on functions;
- if the SYCL backend does not support the generic address space then the function cannot use raw pointers as parameter or return types. Explicit pointer classes must be used instead;
- externally defined functions cannot call a sycl::parallel_for_work_item member function;
- externally defined functions cannot be called from a sycl::parallel_for_work_group scope.

The SYCL linkage mechanism is optional and implementation defined.

6. SYCL Extensions

This section describes the mechanism by which the SYCL core specification can be extended. An extension can be either of two flavors: an extension ratified by the Khronos SYCL group or a vendor supplied extension. In both cases, an extension is an optional feature set which an implementation need not implement in order to be conformant with the core SYCL specification.

Vendors may choose to define extensions in order to expose custom features or to gather feedback on an API that is not yet ready for inclusion in the core specification. Since the APIs for extensions may change as feedback is gathered, the extension mechanism includes a way for application developers to test for the API version of each extension. Once a vendor extension has stabilized, vendors are encouraged to promote it to a future version of the core SYCL specification. Thus, extensions can be viewed as a pipeline of features for consideration in future SYCL versions.

This section does not describe any particular extension to SYCL. Rather, it describes the *mechanism* for defining an extension. Each extension is defined by its own separate document. If an extension is ratified by the Khronos SYCL group, that group will release a document describing the extension. If a vendor defines an extension, the vendor is responsible for releasing its documentation.

6.1 Definition of an extension

An extension can be implemented by adding new types or free functions in a specific namespace, by adding functionality to an existing class that is defined in the core SYCL specification, or through a combination of the two.

New types or free functions for Khronos ratified extensions are defined in the namespace ::sycl::khr::< extensionname>. For example, ::sycl::khr::fancy could be the namespace for a Khronos extension named "fancy".

If a vendor specific extension adds new types or free functions, the vendor is encouraged to define them in the namespace ::sycl::ext::<vendorname> and they are encouraged to add another namespace layer according to the name of the extension. For example, ::sycl::ext::acme::fancier could be the namespace for an extension from the Acme vendor. However, vendors may also choose to define new types and free functions in another top-level namespace that is outside of ::sycl. This might be more appropriate, for example, when an extension integrates features from an existing non-SYCL API. A vendor may not define new types or free functions underneath ::sycl.unless they are in ::sycl::ext::<vendorname>.

[Note: Vendors are discouraged from defining top level namespaces that start with the word "sycl" because we believe that application developers may want to use namespaces like this as namespace aliases. — end note]

Extensions may only add functionality to existing SYCL classes in a limited way. When a Khronos ratified extension needs to add functionality to an existing class, it does so by adding a member function named khr() to that class. For example, an extension on the device class would add a member function like this:

¹ class device {

```
2 // ...
3 sycl::khr::device khr();
4 }:
```

The khr() member function returns an object, and that object provides member functions that are part of the extension.

Likewise, a vendor specific extension may add functionality to an existing SYCL class by adding a member function named ext_<vendorname>() (e.g. ext_acme() for the Acme vendor) like this:

```
1 class device {
2     // ...
3     sycl::ext::acme::device ext_acme();
4  };
```

One motivation for this pattern is to reduce verbosity of application code that uses an extension and to facilitate application migration when an extension is promoted to the core SYCL specification. Consider the following application code:

```
1 void foo(sycl::device dev) {
2  dev.ext_acme().fancy();
3 };
```

If the extension "fancy" is later promoted to the core SYCL specification, the application need only remove the call to ext_acme() in order to migrate the application.

Extensions may also add C++ attributes. The attribute namespace sycl:: is reserved for attributes in the core SYCL specification and for Khronos ratified extensions. Vendor defined extensions should use a different attribute namespace.

Applications must include a special header file in order to get declarations for the types and free functions of an extension. Each Khronos ratified extension has an associated header named "SYCL/khr/<extensionname>.hpp".

The include path "SYCL/ext/<vendorname>" is reserved for vendor extensions. Vendors can choose to provide a single header for all extensions or to provide separate headers for each extension. For example, the Acme vendor could provide the header "SYCL/ext/acme/extensions.hpp" for access to all of its extensions. As with namespaces, vendors are encouraged to define header files in "SYCL/ext/<vendorname>", but a vendor may also define header files in another file system path that is outside of the "SYCL" directory. Vendors may not define header files in the "SYCL" path unless they are underneath "SYCL/ext/<vendorname>".

6.2 Predefined macros

Each Khronos ratified extension has a corresponding feature test macro of the form SYCL_KHR_<extensionname> whose value follows the C++20 pattern for language feature test macros. The value is a number with 6 decimal digits in YYYYMM format identifying the year and month the extension was first adopted or the date the extension was last updated. An implementation must predefine this macro only if it implements the extension, so applications can use the macro in order to determine if the extension is available.

If an implementation provides a vendor specific extension, it should also predefine a feature test macro of the form SYCL_EXT_<vendorname>_<extensionname> (e.g. SYCL_EXT_ACME_FANCY). The value of the macro must be

an integer that monotonically increases for each version of the extension, and vendors are encouraged to use the same YYYYMM format described above.

[Note: The feature test macros are defined uniformly across all parts of a SYCL application, just like any macro. If an implementation uses SMCP, all compiler passes predefine a particular feature test macro the same way, regardless of whether that compiler pass's device supports the feature. Thus, the feature test macros cannot be used to determine whether any particular device supports a feature. If the feature is device-specific, the application must use device::has() or platform::has() to test the feature's aspect in order to determine whether a particular device supports the feature is device.

Each vendor's implementation must also predefine a macro of the form SYCL_VENDOR_<vendorname> (e.g. SYCL_VENDOR_ACME), which applications can use to determine whether they are being compiled by that vendor's toolchain.

An implementation, of course, is allowed to predefine additional macros too. However, an implementation may not predefine a macro whose name starts with SYCL unless it starts with SYCL_EXT_<vendorname> or SYCL_VENDOR_<vendorname>.

6.3 Device aspects and conditional features

An extension may define additional device aspects and it may provide features which are only available on devices with certain aspects. If it does so, the extension documentation must describe which aspects enable these conditional features. If an extension provides a new enumerated aspect value, the type of the new value must be ::sycl::aspect but the enumerated value must be in the extension's namespace scope. For example, a Khronos ratified extension could add a new aspect value like this:

```
1 namespace sycl {
2 namespace khr {
3 namespace aspect {
4
5
    static constexpr auto foo = static_cast<sycl::aspect>(1000);
6
7
    } // namespace aspect
8
    } // namespace khr
9
10
   template<> struct is_aspect_active<khr::aspect::foo> : std::true_type {};
11
12 } // namespace sycl
```

A vendor extension could add an aspect value in a similar way:

```
1 namespace sycl {
2 namespace ext {
3 namespace acme {
4 namespace aspect {
5
6 static constexpr auto bar = static_cast<sycl::aspect>(-1);
7
8 } // namespace aspect
9 } // namespace acme
10 } // namespace ext
```

CHAPTER 6. SYCL EXTENSIONS

```
11
12 template<> struct is_aspect_active<ext::acme::aspect::bar> : std::true_type {};
13
14 } // namespace sycl
```

In the examples above, the vendor has decided to implement aspects from Khronos ratified extensions starting at 1000 and to implement vendor specific aspects as negative integers. However, these are just example implementation details. The SYCL specification does not prescribe the numerical value of any aspect.

6.4 Backends

A vendor extension may define a new SYCL backend. If it does so, the enumerated value for the backend should be defined in the extension's namespace, similar to the way an extended aspect is defined:

```
1 namespace sycl {
2 namespace ext {
3 namespace acme {
4 namespace backend {
5
6 static constexpr auto foo = static_cast<sycl::backend>(-1);
7
8 } // namespace backend
9 } // namespace acme
10 } // namespace ext
11 } // namespace sycl
```

The backend's interoperability API should be made available through a header named "SYCL/ext/<vendorname >/backend/<backendname>.hpp" and it should be defined in the namespace ::sycl::ext::<vendorname>::< backendname>. The implementation should also predefine a macro of the form SYCL_EXT_<vendorname>_BACKEND_ <backendname> when the backend is active.

6.5 Conditional features and compilation errors

SYCL applications are allowed to contain kernels for heterogeneous devices and those kernels, of course, are allowed to use features that are available only on certain devices. Applications are responsible for ensuring that a kernel using such a feature is never submitted to a device that does not support the feature and is never compiled for a device that does not support the feature (e.g. via the module build() or compile() functions). If an application fails to adhere to this requirement, the implementation raises a feature_not_supported exception.

[Note: If an implementation defines a compiler flag that causes some kernels to be pre-compiled for some devices, the vendor is responsible for defining the semantics about when errors are reported for kernels that use device specific extensions. — end note]

An implementation may not raise a spurious error as a result of speculative compilation of a kernel for a device when the application did not specifically ask to submit the kernel to that device or to compile the kernel for that device. To clarify, consider the following example. An application with kernels K1 and K2 runs on devices D1 and D2. Kernel K1 uses extensions specific to D1, and kernel K2 uses extensions specific to D2. The application is coded to ensure that K1 is only submitted to D1 and that K2 is only submitted to D2. An implementation may not raise errors due to speculative compilation of K1 for device D2 or for compilation of K2 for device D1.

An implementation is required, however, to raise an error for a kernel that is not valid for any device. Therefore an implementation must raise an error for a kernel K that is invalid for all devices, even if the application is coded such that kernel K is never submitted to any device.

CHAPTER 6. SYCL EXTENSIONS

A. Information descriptors

The purpose of this chapter is to include all the headers of the memory object descriptors, which are described in detail in Chapter 4, for platform, context, device, and queue.

A.1 Platform information descriptors

The following interface includes all the information descriptors for the platform class as described in Table 4.12.

```
1 namespace sycl {
2 namespace info {
3
   enum class platform : unsigned int {
4
     profile,
5
     version,
6
     name.
7
     vendor,
8
     extensions // Deprecated
9 };
10 } // namespace info
11 } // namespace sycl
```

A.2 Context information descriptors

The following interface includes all the information descriptors for the context class as described in Table 4.15.

```
1 namespace sycl {
2 namespace info {
3 enum class context : int {
4
     platform,
5
     devices,
     atomic_memory_order_capabilities,
6
7
     atomic_fence_order_capabilities
8
     atomic_memory_scope_capabilities,
9
     atomic_fence_scope_capabilities
10 };
11
   } // info
12 } // sycl
```

A.3 Device information descriptors

The following interface includes all the information descriptors for the device class as described in Table 4.19.

1 namespace sycl { 2 namespace info { 3 4 enum class device : int { 5 device_type, 6 vendor_id, 7 max_compute_units, 8 max_work_item_dimensions, 9 max_work_item_sizes, 10 max_work_group_size, 11 preferred_vector_width_char, 12 preferred_vector_width_short, 13 preferred_vector_width_int, 14 preferred_vector_width_long, 15 preferred_vector_width_float, 16 preferred_vector_width_double, 17 preferred_vector_width_half, 18 native_vector_width_char, 19 native_vector_width_short, 20 native_vector_width_int, 21 native_vector_width_long, 22 native_vector_width_float, 23 native_vector_width_double, 24 native_vector_width_half, 25 max_clock_frequency, 26 address_bits, 27 max_mem_alloc_size, 28 image_support, // Deprecated 29 max_read_image_args, 30 max_write_image_args, 31 image2d_max_height, 32 image2d_max_width, 33 image3d_max_height, 34 image3d_max_width, image3d_max_depth, 35 36 image_max_buffer_size, 37 image_max_array_size, 38 max_samplers, 39 max_parameter_size, 40 mem_base_addr_align, 41 half_fp_config, 42 single_fp_config, 43 double_fp_config, 44 global_mem_cache_type, 45 global_mem_cache_line_size, 46 global_mem_cache_size, 47 global_mem_size, 48 max_constant_buffer_size, 49 max_constant_args, 50 local_mem_type, 51 local_mem_size, 52 error_correction_support, 53 host_unified_memory, 54 atomic_memory_order_capabilities,

55 atomic_fence_order_capabilities,

56 atomic_memory_scope_capabilities, 57 atomic_fence_scope_capabilities, 58 profiling_timer_resolution, 59 is_endian_little, 60 is_available, 61 is_compiler_available, // Deprecated 62 is_linker_available, // Deprecated 63 execution_capabilities, 64 queue_profiling, // Deprecated 65 built_in_kernels, 66 platform, 67 name, 68 vendor, 69 driver_version, 70 profile, 71 version, 72 backend_version, 73 aspects, 74 extensions, // Deprecated 75 printf_buffer_size, 76 preferred_interop_user_sync, 77 parent_device, 78 partition_max_sub_devices, 79 partition_properties, 80 partition_affinity_domains, 81 partition_type_property, 82 partition_type_affinity_domain, 83 reference_count 84 }; 85 86 enum class device_type : unsigned int { 87 cpu, // Maps to OpenCL CL_DEVICE_TYPE_CPU 88 gpu, // Maps to OpenCL CL_DEVICE_TYPE_GPU 89 accelerator, // Maps to OpenCL CL_DEVICE_TYPE_ACCELERATOR 90 // Maps to OpenCL CL_DEVICE_TYPE_CUSTOM custom, 91 automatic, // Maps to OpenCL CL_DEVICE_TYPE_DEFAULT 92 host, 93 // Maps to OpenCL CL_DEVICE_TYPE_ALL all 94 }; 95 96 enum class partition_property : int { 97 no_partition, 98 partition_equally, 99 partition_by_counts, 100 partition_by_affinity_domain 101 }; 102 103 enum class partition_affinity_domain : int { 104 not_applicable, 105 numa, 106 L4_cache, 107 L3_cache, 108 L2_cache, 109 L1_cache, 110 next_partitionable

APPENDIX A. INFORMATION DESCRIPTORS

```
111 };
112
113
    enum class local_mem_type : int { none, local, global };
114
115 enum class fp_config : int {
116
      denorm,
117
      inf_nan,
118
      round_to_nearest.
119
      round_to_zero,
120
      round_to_inf,
121
      fma,
122
      correctly_rounded_divide_sqrt,
123
      soft_float
124 };
125
126 enum class global_mem_cache_type : int { none, read_only, read_write };
127
128 enum class execution_capability : unsigned int {
129
      exec_kernel,
130
      exec_native_kernel
131 };
132
133 } // namespace info
134 } // namespace sycl
```

A.4 Queue information descriptors

The following interface includes all the information descriptors for the queue class as described in Table 4.23.

```
1 namespace sycl {
2 namespace info {
3 enum class queue : int {
4 context,
5 device
6 };
7 } // namespace info
8 } // namespace sycl
```

A.5 Kernel information descriptors

The following interface includes all the information descriptors for the kernel class as described in Table 4.98.

```
1 namespace sycl {
2 namespace info {
3 enum class kernel: int {
4 function_name,
5 num_args,
6 context,
7 module,
8 attributes
9 };
```

APPENDIX A. INFORMATION DESCRIPTORS

```
10
11 enum class kernel_work_group: int {
12
    global_work_size,
13
    work_group_size,
14
    compile_work_group_size,
15
    preferred_work_group_size_multiple,
16
    private_mem_size
17 };
18
19 enum class kernel_device_specific: int {
20
    global_work_size,
21
    work_group_size,
22
    compile_work_group_size,
23
    preferred_work_group_size_multiple,
24
    private_mem_size,
25
    max_num_sub_groups,
26
    compile_num_sub_groups,
27
    max_sub_group_size,
28
    compile_sub_group_size
29 };
30
31 } // namespace info
32 } // namespace sycl
```

A.6 Event information descriptors

The following interface includes all the information descriptors for the event class as described in Table 4.28 and Table 4.29.

```
1 namespace sycl {
2 namespace info {
   enum class event: int {
3
4
     command_execution_status
5
   };
6
7
   enum class event_command_status : int {
8
     submitted,
9
     running,
10
     complete
11 };
12
13 enum class event_profiling : int {
14
     command_submit,
15
     command_start,
16
     command_end
17 };
18 } // namespace info
19 } // namespace sycl
```

APPENDIX A. INFORMATION DESCRIPTORS
B. Feature sets

As of SYCL 2020 there are now two distinct feature sets which a SYCL implementation can conform to, in order to better fit the requirements of different domains, such as embedded, mobile, and safety critical, which may have limitations because of the toolchains used.

A SYCL implementation can choose to conform to either the full feature set or the reduced feature set.

B.1 Full feature set

The full feature set includes all features specified in the core SYCL specification with no exceptions.

B.2 Reduced feature set

The reduced feature set makes certain features optional or restricted to specific forms. The following list defines all the differences between the reduced feature set and the full feature set.

1. Un-named SYCL kernel functions: SYCL kernel functions which are are defined using a lambda expression and therefore have no standard name are required to be provided a name via the kernel name template parameter of kernel invocation functions such as parallel_for. This overrides the core SYCL specification rules for SYCL kernel function naming as specified in Section 4.10.7.

B.3 Compatibility

In order to avoid introducing any kind of divergence the reduced and full feature sets are defined such that the full feature set is a subsumption of the reduced feature set. This means that any applications which are developed for the reduced feature will be compatible with both a SYCL reduced implementation and a SYCL full implementation.

B.4 Conformance

One of the reasons for having this be defined in the specification is that hardware vendors which wish to support SYCL on their platform(s) want to be able to demonstrate their support for it by passing conformance. However, if passing conformance means adopting features which they do not believe to be necessary at an additional development effort then this may deter them.

Each feature set has its own route for passing conformance allowing adopters of SYCL to specify the feature set they wish to test conformance against. The conformance test suite would then alter or disable the tests within the test suite according to how the feature sets are differentiated above.

APPENDIX B. FEATURE SETS

C. Host backend specification

This chapter describes how SYCL is mapped on the SYCL host backend. The SYCL host backend exposes the host where the SYCL application is executing as a platform to dispatch SYCL kernels. The SYCL host backend exposes at least one SYCL host device.

C.1 Mapping of the SYCL programming model on the host

The SYCL host device implements all functionality required to execute the SYCL kernels directly on the host, without relying on a third party API. It has full SYCL capabilities and reports them through the SYCL information retrieval interface. At least one SYCL host device must be exposed in the SYCL host backend in all SYCL implementations, and it must always be available. Any C++ application debugger, if available on the system, can be used for debugging SYCL kernels executing on a SYCL host device.

When a SYCL implementation executes kernels on the host device, it is free to use whatever parallel execution facilities available on the host, as long as it executes within the semantics of the kernel execution model defined by the SYCL kernel execution model.

Kernel math library functions on the host must conform to OpenCL math precision requirements. The SYCL host device needs to be queried for the capabilities it provides. This ensures consistency when executing any SYCL general application.

The SYCL host device must report as supporting images and therefore support the minimum image formats.

The range of image formats supported by the host device is implementation-defined, but must match the minimum requirements of the OpenCL specification.

SYCL implementors can provide extensions on the host-device to match any other backend-specific extension. This allows developers to rely on the host device to execute their programs when said backend is not available.

C.1.1 SYCL memory model on the host

All SYCL device memories are available on devices from the host backend.

| SYCL | Host |
|----------|---------------|
| Global | System memory |
| Constant | System memory |
| Local | System memory |
| Private | Stack |

Table C.1: Mapping of SYCL memory regions into host memory regions.

C.2 Interoperability with the host application

The host backend must ensure all functionality of the SYCL generic programming model is always available to developers. However, since there is no heterogeneous API behind the host backend (it directly targets the host platform), there are no native types for SYCL objects to map to in the SYCL application.

Inside SYCL kernels, the host backend must ensure interoperability with existing host code, so that existing host libraries can be used inside SYCL kernels executing on the host. In particular, when retrieving a raw pointer from a multi pointer object, the pointer returned must be usable by any library accessible by the SYCL application.

D. OpenCL backend specification

This chapter describes how the SYCL general programming model is mapped on top of OpenCL, and how the SYCL generic interoperability interface must be implemented by vendors providing SYCL for OpenCL implementations to ensure SYCL applications written for the OpenCL backend are interoperable.

D.1 SYCL for OpenCL framework

The SYCL framework allows applications to use a host and one or more OpenCL devices as a single heterogeneous parallel computer system. The framework contains the following components:

- SYCL C++ template library: The template library provides a set of C++ templates and classes which provide the programming model to the user. It enables the creation of runtime classes such as SYCL queues, buffers and images, as well as access to some underlying OpenCL runtime object, such as contexts, platforms, devices and program objects.
- SYCL runtime: The SYCL runtime interfaces with the underlying OpenCL implementations and handles scheduling of commands in queues, moving of data between host and devices, manages contexts, programs, kernel compilation and memory management.
- *OpenCL Implementation(s)*: The SYCL system assumes the existence of one or more OpenCL implementations available on the host machine. If no OpenCL implementation is available, then the SYCL implementation provides only the SYCL host device to run kernels on.
- SYCL device compilers: The SYCL device compilers compile SYCL C++ kernels into a format which can be executed on an OpenCL device at runtime. There may be more than one SYCL device compiler in a SYCL implementation. The format of the compiled SYCL kernels is not defined. A SYCL device compiler may, or may not, also compile the host parts of the program.

The OpenCL backend is enabled using the sycl::backend::opencl value of enum class backend. That means that when the OpenCL backend is active, the value of sycl::is_backend_active<sycl::backend::opencl>:: value will be true.

D.2 Mapping of SYCL programming model on top of OpenCL

The SYCL programming model was originally designed as a high-level model for the OpenCL API, hence the mapping of SYCL on the OpenCL API is mostly straightforward.

When the OpenCL backend is active on a SYCL application, all visible OpenCL platforms are exported as SYCL platforms.

When a SYCL implementation executes kernels on an OpenCL device, it achieves this by enqueuing OpenCL **commands** to execute computations on the processing elements within a device. The processing elements within an OpenCL compute unit may execute a single stream of instructions as ALUs within a SIMD unit (which execute

| SYCL | OpenCL |
|----------|-----------------|
| Global | Global memory |
| Constant | Constant memory |
| Local | Local memory |
| Private | Private memory |

Table D.1: Mapping of SYCL memory regions into OpenCL memory regions.

in lockstep with a single stream of instructions), as independent SPMD units (where each PE maintains its own program counter) or as some combination of the two.

D.2.1 Platform mixed version support

The SYCL system presents the user with a set of devices, grouped into some number of platforms. The device version is an indication of the device's capabilities, as represented by the device information returned by the sycl ::device::get_info() member function. Examples of attributes associated with the device version are resource limits and information about functionality beyond the core SYCL specification's requirements. The version returned corresponds to the highest version of the OpenCL specification for which the device is conformant, but is not higher than the version of the device's platform which bounds the overall capabilities of the runtime operating the device.

D.2.2 OpenCL memory model

The memory model for SYCL devices running on OpenCL platforms follows the memory model of the OpenCL version they conform to. Work-items executing in a kernel have access to four distinct memory regions, with the mapping between SYCL and OpenCL described in table D.1.

D.2.3 OpenCL resources managed by SYCL application

In OpenCL, a developer must create a context to be able to execute commands on a device. Creating a context involves choosing a platform and a list of devices. In SYCL, contexts, platforms and devices all exist, but the user can choose whether to specify them or have the SYCL implementation create them automatically. The minimum required object for submitting work to devices in SYCL is the queue, which contains references to a platform, device and context internally.

The resources managed by SYCL are:

- 1. Platforms: all features of OpenCL are implemented by platforms. A platform can be viewed as a given hardware vendor's runtime and the devices accessible through it. Some devices will only be accessible to one vendor's runtime and hence multiple platforms may be present. SYCL manages the different platforms for the user. In SYCL, a platform resource is accessible through a sycl::platform object. SYCL also provides a host platform object, which only contains a single host device.
- 2. Contexts: any OpenCL resource that is acquired by the user is attached to a context. A context contains a collection of devices that the host can use and manages memory objects that can be shared between the devices. Data movement between devices within a context may be efficient and hidden by the underlying OpenCL runtime while data movement between contexts may involve the host. A given context can only wrap devices owned by a single platform. In SYCL, a context resource is accessible through a sycl:: context object.

- 3. Devices: platforms provide one or more devices for executing kernels. In SYCL, a device is accessible through a sycl::device object.
- 4. Kernels: the SYCL functions that run on SYCL devices (i.e. either an OpenCL device, or the host device) are defined as C++ function objects (a named function object type or a lambda function).
- 5. Modules: OpenCL objects that store implementation data for the SYCL kernels. These objects are only required for advanced use in SYCL and are encapsulated in the sycl::module class.
- 6. Queues: SYCL kernels execute in command queues. The user must create a queue, which references an associated context, platform and device. The context, platform and device may be chosen automatically, or specified by the user. In SYCL, command queues are accessible through sycl::queue objects.

D.3 Interoperability with the OpenCL API

The OpenCL backend for SYCL ensures maximum compatibility between SYCL and OpenCL kernels and API. This includes supporting devices with different capabilities and support for different versions of the OpenCL C language, in addition to supporting SYCL kernels written in C++.

SYCL runtime classes which encapsulate an OpenCL opaque type such as SYCL context or SYCL queue must provide an interoperability constructor taking an instance of the OpenCL opaque type. These constructors must retain that instance to increase the reference count of the OpenCL resource.

The destructor for the SYCL runtime classes which encapsulate an OpenCL opaque type must release that instance to decrease the reference count of the OpenCL resource.

Note that an instance of a SYCL runtime class which encapsulates an OpenCL opaque type can encapsulate any number of instances of the OpenCL type, unless it was constructed via the interoperability constructor, in which case it can encapsulate only a single instance of the OpenCL type.

The lifetime of a SYCL runtime class that encapsulates an OpenCL opaque type and the instance of that opaque type retrieved via the get() member function are not tied in either direction given correct usage of OpenCL reference counting. For example if a user were to retrieve a cl_command_queue instance from a SYCL queue instance and then immediately destroy the SYCL queue instance, the cl_command_queue instance is still valid. Or if a user were to construct a SYCL queue instance from a cl_command_queue instance and then immediately release the cl_command_queue instance, the SYCL queue instance is still valid.

Note that a SYCL runtime class that encapsulates an OpenCL opaque type is not responsible for any incorrect use of OpenCL reference counting outside of the SYCL runtime. For example if a user were to retrieve a cl_command_queue instance from a SYCL queue instance and then release the cl_command_queue instance more than once without any prior retain then the SYCL queue instance that the cl_command_queue instance was retrieved from is now undefined.

Note that an instance of the SYCL buffer or SYCL image class templates constructed via the interoperability constructor is free to copy from the cl_mem into another memory allocation within the SYCL runtime to achieve normal SYCL semantics, for as long as the SYCL buffer or SYCL image instance is alive.

Table D.2 relates SYCL objects to their OpenCL native type in the SYCL application.

| OpenCL backend native types | Description | |
|---|--|--|
| device | A SYCL device object encapsulates an | |
| cl_device_id | OpenCL device object. | |
| context | A SYCL context object encapsulates an | |
| cl_context | OpenCL context object. | |
| program | When a SYCL program is constructed for | |
| cl_program | the OpenCL backend, this maps directly to | |
| | an OpenCL program object. | |
| kernel | The SYCL implementation will produce | |
| cl_kernel | OpenCL programs from the SYCL device | |
| | kernels. They are dispatched on the OpenCL | |
| | interface as OpenCL kernel objects. This | |
| | also apply to built-in kernels. | |
| event | A SYCL event can encapsulate one or mul- | |
| <pre>std::vector<cl_event></cl_event></pre> | tiple OpenCL events, representing a number | |
| | of dependencies in the same or different con- | |
| | texts, that must be satisfied for the SYCL | |
| | event to be complete. | |
| buffer | SYCL buffers containing OpenCL memory | |
| <pre>std::vector<cl_mem></cl_mem></pre> | objects can handle multiple cl_mem objects | |
| | in the same or different context. The inter- | |
| | operability interface will return a list of ac- | |
| | tive buffers in the SYCL runtime. | |
| image | SYCL images containing OpenCL image | |
| <pre>std::vector<cl_mem></cl_mem></pre> | objects can handle multiple underlying | |
| | cl_mem objects at the same time in the same | |
| | or different OpenCL contexts. The interop- | |
| | erability interface will return a list of active | |
| | images in the SYCL runtime. | |
| | End of table | |

Table D.2: List of native types per SYCL object in the OpenCL backend.

The user can also extract OpenCL cl_kernel and cl_program objects for kernels by providing the type name of the kernel.

Inside the SYCL kernel, the SYCL API offers interoperability with OpenCL device types. The table D.3 describes the mapping of kernel types.

| SYCL kernel native types in OpenCL | Description |
|------------------------------------|--|
| <pre>multi_ptr::get()</pre> | Returns a pointer in the OpenCL address |
| | space corresponding to the type of multi |
| | pointer object |
| <pre>device_event::get()</pre> | Returns an event_t object, which can be |
| | used to identify copies from global to local |
| | memory and vice-versa |
| | End of table |

Table D.3: List of native types per SYCL object on kernel code.

When a buffer or image is allocated on more than one OpenCL device, if these devices are on separate contexts then multiple cl_mem objects may be allocated for the memory object, depending on whether the object has actively been used on these devices yet or not.

Some types in SYCL vary according to pointer size or vary on the host according to the host ABI, such as size_t or long. In order for the the SYCL device compiler to ensure that the sizes of these types match the sizes on the host and to enable data of these types to be shared between host and device, the OpenCL interoperability types are defined, sycl::cl_int and sycl::cl_size_t.

The OpenCL C function qualifier __kernel and the access qualifiers: __read_only, __write_only and __read_write are not exposed in SYCL via keywords, but are instead encapsulated in SYCL's parameter passing system inside accessors. Users wishing to achieve the OpenCL equivalent of these qualifiers in SYCL should instead use SYCL accessors with equivalent semantics.

Any OpenCL C function included in a pre-built OpenCL library can be defined as an extern "C" function and the OpenCL program has to be linked against any SYCL program that contains kernels using the external function. In this case, the data types used have to comply with the interoperability aliases defined in D.7.

D.4 Programming interface

The following section describes the OpenCL-specific API. All free functions are available in the sycl::opencl namespace.

| OpenCL interoperability functions | Description | |
|--|--|--|
| <pre>sycl::context make_context (</pre> | Constructs a SYCL context instance from | |
| <pre>const cl_context &clContext,</pre> | an OpenCL cl_context in accordance with | |
| <pre>const sycl::async_handler &asyncHandler = {})</pre> | the requirements described in 4.5.2. | |
| <pre>cl_context get_native<sycl::context>(</sycl::context></pre> | Returns a valid cl_context instance in ac- | |
| <pre>const sycl::context &syclContext)</pre> | cordance with the requirements described in | |
| | 4.5.2. | |
| <pre>sycl::event make_event (const cl_event &clEvent,</pre> | Constructs a SYCL event instance from | |
| <pre>const sycl::context &syclContext)</pre> | an OpenCL cl_event in accordance with | |
| | the requirements described in 4.5.2. The | |
| | syclContext must match the OpenCL con- | |
| | text associated with the clEvent. | |
| <pre>cl_event get_native<sycl::event>(</sycl::event></pre> | Returns a valid cl_event instance in accor- | |
| <pre>const sycl::event &syclEvent)</pre> | dance with the requirements described in | |
| | 4.5.2. | |
| <pre>sycl::device make_device(</pre> | Constructs a SYCL device instance from an | |
| <pre>const cl_device_id &clDeviceId)</pre> | OpenCL cl_device_id in accordance with | |
| | the requirements described in 4.5.2. | |
| <pre>cl_device_id get_native<sycl::device>(</sycl::device></pre> | Returns a valid cl_device_id instance in ac- | |
| <pre>const sycl::device &syclDevice)</pre> | cordance with the requirements described in | |
| | 4.5.2. | |
| <pre>sycl::platform make_platform(</pre> | Constructs a SYCL platform instance from | |
| <pre>const cl_platform_id &clPlatformId)</pre> | an OpenCL cl_platform_id in accordance | |
| | with the requirements described in 4.5.2. | |
| <pre>cl_platform_id get_native<sycl::platform>(</sycl::platform></pre> | Returns a valid cl_platform_id instance in | |
| <pre>const sycl::platform &syclPlatform)</pre> | accordance with the requirements described | |
| | in 4.5.2. | |

| OpenCL interoperability functions | Description |
|--|---|
| <pre>sycl::queue make_queue(</pre> | Constructs a SYCL queue instance with an |
| <pre>const cl_command_queue &clQueue,</pre> | optional async_handler from an OpenCL |
| <pre>const sycl::context &syclContext,</pre> | cl_command_queue in accordance with the |
| <pre>const sycl::async_handler &asyncHandler = {})</pre> | requirements described in 4.5.2. |
| <pre>cl_command_queue get_native<sycl::queue>(</sycl::queue></pre> | Returns a valid cl_command_queue instance |
| <pre>const sycl::queue &syclQueue)</pre> | in accordance with the requirements de- |
| | scribed in 4.5.2. |
| <pre>sycl::buffer make_buffer(</pre> | Available only when: dimensions == 1. |
| <pre>const cl_mem &clMemObject,</pre> | Constructs a SYCL buffer instance from |
| <pre>const sycl::context &syclContext,</pre> | an OpenCL cl_mem in accordance with the |
| <pre>sycl::event availableEvent = {})</pre> | requirements described in 4.5.2. The in- |
| | stance of the SYCL buffer class tem- |
| | plate being constructed must wait for the |
| | SYCL event parameter, if one is provided, |
| | availableEvent to signal that the cl_mem |
| | instance is ready to be used. The SYCL |
| | context parameter syclContext is the con- |
| | text associated with the memory object. |
| <pre>sycl::sampled_image make_sampled_image(</pre> | Constructs a SYCL sampled_image instance |
| <pre>const cl_mem &clMemObject,</pre> | from an OpenCL cl_mem in accordance with |
| <pre>const context &syclContext,</pre> | the requirements described in 4.5.2. The |
| <pre>event availableEvent = {})</pre> | instance of the SYCL image class tem- |
| | plate being constructed must wait for the |
| | SYCL event parameter, if one is provided, |
| | availableEvent to signal that the cl_mem |
| | instance is ready to be used. The SYCL |
| | context parameter syclContext is the con- |
| | text associated with the memory object. |
| syc1::unsampled_image make_unsampled_image(| stones from on OpenCL al men in secon |
| const ci_mem &cimemobject, | dance with the requirements described in |
| const syc1::context asyc1context, | 4.5.2 The instance of the SVCL image class |
| sycimage_sampler syciimagesampler, | 4.5.2. The instance of the STCL image class |
| event availableEvent = {}) | SVCL event parameter if one is provided |
| | available Event to signal that the cl more |
| | instance is ready to be used. The SVCI |
| | context parameter syclContext is the con- |
| | text associated with the memory object |
| svclimage sampler make image sampler(| Constructs a SYCL image sampler instance |
| const cl sampler &clSampler | from an OpenCL cl sampler in accordance |
| const svcl::context &svclContext | with the requirements described in 4.5.2 |
| const sycicontext asycicontext, | The SYCL context parameter sycl Context |
| | is the context associated with the sampler |
| | object. |

| OpenCL interoperability functions | Description | |
|---|---|--|
| <pre>sycl::kernel make_kernel(</pre> | Constructs a SYCL kernel instance from | |
| <pre>const cl_kernel &clKernel,</pre> | an OpenCL cl_kernel in accordance with | |
| <pre>const sycl::context& syclContext)</pre> | the requirements described in 4.5.2. The | |
| | SYCL context must represent the same un- | |
| | derlying OpenCL context associated with | |
| | the OpenCL kernel object. | |
| <pre>cl_kernel get_native<sycl::kernel>(</sycl::kernel></pre> | Returns a valid cl_kernel instance in ac- | |
| <pre>const sycl::kernel &syclKernel)const</pre> | cordance with the requirements described in | |
| | 4.5.2. | |

D.4.1 Reference counting

All OpenCL objects are reference counted. The SYCL general programming model doesn't require that native objects are reference counted. However, for convenience, the following function is provided in the sycl::opencl namespace.

| Reference counting | Description |
|--|---|
| <pre>template <typename openclt=""></typename></pre> | Returns the reference counting of the given |
| <pre>cl_uint get_reference_count(openCLT obj)</pre> | object |

D.4.2 Errors and limitations

If there is an OpenCL error associated with an exception triggered, then the OpenCL error code can be obtained by the free function cl_int sycl::opencl::get_error_code(sycl::exception&). In the case where there is no OpenCL error associated with the exception triggered, the OpenCL error code will be CL_SUCCESS.

D.4.3 Interoperability with modules

In OpenCL [1] any kernel function that is enqueued over an nd-range is represented by acl_kernel and must be compiled and linked via a cl_program using clBuildProgram, clCompileProgram and clLinkProgram.

For OpenCL SYCL backend this detail is abstracted away by modules and a module object containing all SYCL kernel functions in a translation is retrieved by calling the free function this_module::get.

However, there are cases where it is useful to be able to manually create a module from an input specific to the OpenCL SYCL backend such as OpenCL C source, and intermediate representation/language such as SPIR-V. This can be useful for interoperability with existing OpenCL kernels or libraries or binaries generated by another tool which need to be linked at runtime.

The OpenCL SYCL backend specification provides additional free functions which provide the above functionality, each resulting in an input module which can then be built, compiled and linked as described in 4.13.7.

```
1 namespace sycl {
2 namespace opencl {
3
4 using binary_blob_t = std::pair<const char*, size_t>;
5
6 module<module_state::input> create_module_with_source (context ctx, std::string source);
```

```
7
8 module<module_state::input> create_module_with_binary(context ctx, binary_blob_t binary);
9
10 module<module_state::input> create_module_with_il (context ctx, binary_blob_t il);
11
12 module<module_state::input> create_module_with_builtin_kernels (context ctx,
13 std::vector<std::string> kernelNames);
14
15 } // namespace opencl
16 } // namespace sycl
```

D.4.3.1 Free functions

- 1 module_state::input> create_module_with_source (context ctx, std::string source); // (1)
 - 1. *Preconditions:* The context specified by ctx must be associated with the OpenCL SYCL backend. The OpenCL C source specified by source must not be an empty string.

Effects: Constructs a module from the provided OpenCL C source specified by source and associated with the context specified by ctx by invoking the necessary OpenCL APIs.

Returns: A module of module_state::input containing the kernels defined in the OpenCL C source specified by source.

Throws: invalid_object_error if any error is produced by invoking the OpenCL APIs.

- 1 module<module_state::input> create_module_with_binary(context ctx, binary_blob_t binary); // (1)
 - 1. *Preconditions:* The context specified by ctx must be associated with the OpenCL SYCL backend. The binary blob specified by binary must not contain a null pointer or zero size.

Effects: Constructs a module from the provided binary blob specified by binary and associated with the context specified by ctx by invoking the necessary OpenCL APIs.

Returns: A module of module_state::input containing the kernels defined in the binary blob specified by binary.

Throws: invalid_object_error if any error is produced by invoking the OpenCL APIs.

- 1 module<module_state::input> create_module_with_il (context ctx, binary_blob_t il); // (1)
 - 1. *Preconditions:* The context specified by ctx must be associated with the OpenCL SYCL backend. The intermediate language specified by i1 must not contain a null pointer or zero size.

Effects: Constructs a module from the provided intermediate language specified by il and associated with the context specified by ctx by invoking the necessary OpenCL APIs.

Returns: A module of module_state::input containing the kernels defined in the binary intermediate language by i1.

Throws: invalid_object_error if any error is produced by invoking the OpenCL APIs.

- 1 module_module_state::input> create_module_with_builtin_kernels (context ctx, //(1)
- 2 std::vector<std::string> kernelNames);
 - 1. *Preconditions:* The context specified by ctx must be associated with the OpenCL SYCL backend. The list of names specified by kernelNames must not be empty.

Effects: Constructs a module from the provided builtin kernel names specified by kernelNames and associated with the context specified by ctx by invoking the necessary OpenCL APIs.

Returns: A module of module_state::input containing the built-in kernels defined by the list of kernel names specified by kernelNames.

Throws: invalid_object_error if any error is produced by invoking the OpenCL APIs.

D.4.4 Interoperability with kernels

It is possible to construct a kernel from a previously created OpenCL cl_kernel by calling the interop free function make_kernel defined in 4.5.2.3.

This will create a kernel object which can be invoked by any of kernel invocation commands such as parallel_for which take a kernel but not SYCL kernel function.

Calling make_kernel must trigger a call to clRetainKernel and the resulting kernel object must call clReleaseKernel on destruction.

The kernel arguments for the OpenCL C kernel kernel can either be set prior to creating the kernel object or by calling the set_arg member function of the handler class.

If kernel arguments are set prior to creating the kernel object the SYCL runtime is not responsible for managing the data of these arguments.

D.4.5 OpenCL kernel conventions and SYCL

OpenCL and SYCL use opposite conventions for the unit stride dimension. SYCL aligns with C++ conventions, which is important to understand from a performance perspective when porting code to SYCL. The unit stride dimension, at least for data, is implicit in the linearization equations in SYCL (Equation 4.3) and OpenCL. SYCL aligns with C++ array subscript ordering arr[a][b][c], in that range constructor dimension ordering used to launch a kernel (e.g. range<3> R{a,b,c}) and range and ID queries within a kernel, are ordered in the same way as the C++ multi-dimensional subscript operators (unit stride on the right).

When specifying a range as the global or local size in a parallel_for that invokes an OpenCL interop kernel (through cl_kernel interop or compile_with_source/build_with_source), the highest dimension of the range in SYCL will map to the lowest dimension within the OpenCL kernel. That statement applies to both an underlying enqueue operation such as clEnqueueNDRangeKernel in OpenCL, and also ID and size queries within the OpenCL kernel. For example, a 3D global range specified in SYCL as:

range<3> R{r0,r1,r2};

maps to an clEnqueueNDRangeKernel global_work_size argument of:

size_t cl_interop_range[3] = {r2,r1,r0};

Likewise, a 2D global range specified in SYCL as:

range<2> R{r0,r1};

maps to an clEnqueueNDRangeKernel global_work_size argument of:

size_t cl_interop_range[2] = {r1,r0};

The mapping of highest dimension in SYCL to lowest dimension in OpenCL applies to all operations where a multi-dimensional construct must be mapped, such as when mapping SYCL explicit memory operations to OpenCL APIs like clEnqueueCopyBufferRect.

Work-item and work-group ID and range queries have the same reversed convention for unit stride dimension between SYCL and OpenCL. For example, with three, two, or one dimensional SYCL global ranges, OpenCL and SYCL kernel code queries relate to the range as shown in Table D.6. The "SYCL kernel query" column applies for SYCL-defined kernels, and the "OpenCL kernel query" column applies for kernels defined through OpenCL interop.

| SYCL kernel query | OpenCL kernel query | Returned Value |
|---|-------------------------------|---------------------------------|
| With enqueued 3D SYCL global range of range<3> R{r0,r1,r2} | | |
| <pre>nd_item::get_global_range(0)/ item::get_range(0) get_global_siz</pre> | | rØ |
| <pre>nd_item::get_global_range(1)/ item::get_range(1)</pre> | <pre>get_global_size(1)</pre> | r1 |
| <pre>nd_item::get_global_range(2)/ item::get_range(2)</pre> | <pre>get_global_size(0)</pre> | r2 |
| <pre>nd_item::get_global_id(0)/ item::get_id(0)</pre> | <pre>get_global_id(2)</pre> | Value in range 0(r0 -1) |
| <pre>nd_item::get_global_id(1)/ item::get_id(1)</pre> | <pre>get_global_id(1)</pre> | Value in range 0(r1-1) |
| <pre>nd_item::get_global_id(2)/ item::get_id(2)</pre> | <pre>get_global_id(0)</pre> | Value in range 0(r2-1) |
| With enqueued 2D SYCL global range of range<2> R{r0,r1} | | |
| <pre>nd_item::get_global_range(0)/ item::get_range(0) get_global_size(1) r0</pre> | | rØ |
| <pre>nd_item::get_global_range(1)/ item::get_range(1)</pre> | <pre>get_global_size(0)</pre> | r1 |
| <pre>nd_item::get_global_id(0)/ item::get_id(0)</pre> | <pre>get_global_id(1)</pre> | Value in range 0(r0-1) |
| <pre>nd_item::get_global_id(1)/ item::get_id(1)</pre> | <pre>get_global_id(0)</pre> | Value in range 0(r1-1) |
| With enqueued 1D SYCL global range of range<1> R{r0} | | |
| <pre>nd_item::get_global_range(0)/ item::get_range(0)</pre> | <pre>get_global_size(0)</pre> | rØ |
| <pre>nd_item::get_global_id(0)/ item::get_id(0)</pre> | <pre>get_global_id(0)</pre> | Value in range 0(r0-1) |

Table D.6: Example range mapping from SYCL enqueued three dimensional global range to OpenCL and SYCL queries.

D.4.6 Data types

The OpenCL C language standard [1, Section 6.11] defines its own built-in scalar data types, and these have additional requirements in terms of size and signedness on top of what is guaranteed by ISO C++. For the purpose of interoperability and portability, SYCL defines a set of aliases to C++ types within the sycl::opencl namespace using the c1_ prefix. These aliases are described in Table D.7

| Scalar data type alias | Description | |
|------------------------|--|--|
| cl_bool | Alias to a conditional data type which can | |
| | be either true or false. The value true ex- | |
| | pands to the integer constant 1 and the value | |
| | false expands to the integer constant 0. | |
| cl_char | Alias to a signed 8-bit integer, as defined by | |
| | the C++ core language. | |
| cl_uchar | Alias to an unsigned 8-bit integer, as defined | |
| | by the C++ core language. | |
| cl_short | Alias to a signed 16-bit integer, as defined | |
| | by the C++ core language. | |
| cl_ushort | Alias to an unsigned 16-bit integer, as de- | |
| | fined by the C++ core language. | |
| cl_int | Alias to a signed 32-bit integer, as defined | |
| | by the C++ core language. | |
| cl_uint | Alias to an unsigned 32-bit integer, as de- | |
| | fined by the C++ core language. | |
| cl_long | Alias to a signed 64-bit integer, as defined | |
| | by the C++ core language. | |
| cl_ulong | Alias to an unsigned 64-bit integer, as de- | |
| | fined by the C++ core language. | |
| cl_float | Alias to a 32-bit floating-point. The float | |
| | data type must conform to the IEEE 754 sin- | |
| | gle precision storage format. | |
| cl_double | Alias to a 64-bit floating-point. The dou- | |
| | ble data type must conform to the IEEE 754 | |
| | double precision storage format. | |
| cl_half | Alias to a 16-bit floating-point. The | |
| | half data type must conform to the | |
| | IEEE 754-2008 half precision storage for- | |
| | mat. An exception with the errc:: | |
| | feature_not_supported error code must be | |
| | thrown if the half type is used in a SYCL | |
| | kernel function which executes on a SYCL | |
| | device that does not support the extension | |
| | khr_fp16. | |
| | End of table | |

Table D.7: Scalar data type aliases supported by SYCL OpenCL backend.

D.5 Preprocessor directives and macros

• SYCL_BACKEND_OPENCL substitutes to one if the OpenCL SYCL backend is active while building the SYCL application.

D.5.1 Offline linking with OpenCL C libraries

SYCL supports linking SYCL kernel functions with OpenCL C libraries during offline compilation or during online compilation by the SYCL runtime within a SYCL application.

Linking with OpenCL C kernel functions offline is an optional feature and is unspecified. Linking with OpenCL C kernel functions online is performed by using the SYCL module class to compile and link an OpenCL C source; using the compile_with_source or build_with_source member functions.

OpenCL C functions that are linked with, using either offline or online compilation, must be declared as extern "C" function declarations. The function parameters of these function declarations must be defined as the OpenCL C interoperability aliases; pointer of the multi_ptr class template, vector_t of the vec class template and scalar data type aliases described in Table D.7.

For example:

```
1 extern "C" typename sycl::decorated_global_ptr<std::int32_t>::pointer my_func(
```

```
2 sycl::float4::vector_t x, double y);
```

D.6 SYCL support of non-core OpenCL features

In addition to the OpenCL core features, SYCL also provides support for OpenCL extensions which provide features in OpenCL via khr extensions.

Some extensions are natively supported within the SYCL interface, however some can only be used via the OpenCL interoperability interface. The SYCL interface required for native extensions must be available. However if the respective extension is not supported by the executing SYCL device, the SYCL runtime must throw an exception with the errc::feature_not_supported error code.

The OpenCL backend exposes khr extensions to SYCL applications through the sycl::aspect enumerated type. Therefore, applications can query for the existence of khr extensions by calling the device::has() or platform ::has() member functions.

All OpenCL extensions are available through the OpenCL interoperability interface, but some can also be used through core SYCL APIs. Table D.8 shows which these are. Table D.8 also shows the mapping from each OpenCL extension name to its associated SYCL device aspect. Note that some aspects are part of the core SYCL specification, and these are in namespace ::sycl::aspect. Other aspects are specific to the OpenCL backend, and these are in namespace ::sycl::aspect.

| SYCL Aspect | OpenCL Extension | Core SYCL API |
|---|--------------------------------------|---------------|
| <pre>aspect::int64_base_atomics</pre> | <pre>cl_khr_int64_base_atomics</pre> | Yes |
| <pre>aspect::int64_extended_atomics</pre> | cl_khr_int64_extended_atomics | Yes |
| aspect::fp16 | cl_khr_fp16 | Yes |
| <pre>opencl::aspect::3d_image_writes</pre> | <pre>cl_khr_3d_image_writes</pre> | Yes |
| opencl::aspect::khr_gl_sharing | cl_khr_gl_sharing | No |
| <pre>opencl::aspect::apple_gl_sharing</pre> | cl_apple_gl_sharing | No |
| opencl::aspect::d3d10_sharing | cl_khr_d3d10_sharing | No |
| opencl::aspect::d3d11_sharing | cl_khr_d3d11_sharing | No |
| opencl::aspect::dx9_media_sharing | cl_khr_dx9_media_sharing | No |
| | | End of table |

Table D.8: SYCL support for OpenCL 1.2 extensions.

D.6.1 Half precision floating-point

The half scalar data type: half and the half vector data types: half1, half2, half3, half4, half8 and half16 must be available at compile-time. However if any of the above types are used in a SYCL kernel function, executing on a device which does not support the extension khr_fp16, the SYCL runtime must throw an exception with the errc::feature_not_supported error code.

The conversion rules for half precision types follow the same rules as in the OpenCL 1.2 extensions specification [5, par. 9.5.1].

The math functions for half precision types follow the same rules as in the OpenCL 1.2 extensions specification [5, par. 9.5.2, 9.5.3, 9.5.4, 9.5.5]. The allowed error in ULP(Unit in the Last Place) is less than 8192, corresponding to Table 6.9 of the OpenCL 1.2 specification [1].

D.6.2 Writing to 3D image memory objects

The accessor class for target access::target::image in SYCL support member functions for writing 3D image memory objects, but this functionality is *only allowed* on a device if the extension cl_khr_3d_image_writes is supported on that device.

D.6.3 Interoperability with OpenGL

Interoperability between SYCL and OpenGL is not directly provided by the SYCL interface, however can be achieved via the SYCL OpenCL interoperability interface.

E. What has changed from previous versions

E.1 What has changed from SYCL 1.2.1 to SYCL 2020

The SYCL runtime moved from namespace cl::sycl provided by #include "CL/sycl.hpp" to namespace sycl provided by #include "SYCL/sycl.hpp" as explained in Section 4.3. The old header file is still available for compatibility with SYCL 1.2.1 applications.

The SYCL specification is now based on the core language of C++17, as described in Section 3.8.1. Features of C++17 are now enabled within the specification, such as deduction guides for class template argument deduction.

Naming of lambda functions passed to kernel invocations is now optional.

The accessor class definition and interface have been significantly modified to simplify commonly written SYCL code. Default modes have been defined and modifiers use tag arguments. Host accessors have now their own class host_accessor as explained in Section 4.7.6.10.

Kernel attributes have been better described and are now applied to the function type of a kernel function, which allows them to be applied directly to lambdas. This means that propagation of the attribute from a function to the calling kernel is no longer required, and attributes are instead applied directly to the kernel function that they impact.

The list of built-in integer math functions was extended with ctz() in Tables 4.155. Specification of clz() was extended with the case of argument is 0.

The classes vector_class, string_class, function_class, mutex_class, shared_ptr_class, weak_ptr_class, hash_class and exception_ptr_class have been removed from the API and the standard classes std:: vector, std::string, std::function, std::mutex, std::shared_ptr, std::weak_ptr, std::hash and std:: exception_ptr are used instead.

operator[] of SYCL accessor for SYCL buffer was changed to return const reference when accessMode ==
access::mode::read.

The specific sycl::buffer API taking std::unique_ptr has been removed. The behavior is the same as in SYCL 1.2.1 but with a simplified API. Since there is still the API taking std::shared_ptr and there is an implicit conversion from a std::unique_ptr prvalue to a std::shared_ptr, the API can still be used as before with a std::unique_ptr to give away memory ownership.

Unified Shared Memory (USM), in Section 4.8, has been added as a pointer-based strategy for data management. It defines several types of allocations with various accessibility rules for host and devices. USM is meant to complement buffers, not replace them.

The queue class received a new property that requires in-order semantics for a queue where operations are executed in the order in which they are submitted.

The queue class received several new member functions to define kernels directly on a queue objects instead of

inside a command group handler in the submit member function.

The program class has been replaced with the module which provides a type-safe and thread-safe interface for compiling and linking SYCL kernel function. The previous member functions of the program class are now free functions. Modules are now retrieved from the this_module::get function which produces a module containing the SYCL kernel functions of the current translation unit.

The module class now supports specialization constants which allow SYCL kernel functions to define constant variables, whose value is not known until the containing module is compiled. The specialization_constant class has been introduced to represent a reference to a specialization constant both in the SYCL application for setting the value and in a SYCL kernel function for retrieving the value.

The kernel_handler class has been introduced as an optional parameter to kernel invocation commands to provide functionality and queries relating to a SYCL kernel function at the kernel scope, including getting the value of a specialization constant.

The constructors for SYCL context and queue are made explicit to prevent ambiguities in the selected constructor resulting from implicit type conversion.

The requirement for C++ standard layout for data shared between host and devices has been softened and now only C++ trivially copyable is required, as explained mainly in Section 4.14.4.

The concept of a group of work-items was generalized to include work-groups and sub-groups. A work-group is represented by the sycl::group class as in SYCL 1.2.1, and a sub-group is represented by the new sycl:: sub_group class.

The host_task member function for the queue has been introduced for en-queueing host tasks on a queue to schedule the SYCL runtime to invoke native C++ functions, conforming to the SYCL memory model. Host tasks also support interoperability with the native SYCL backend objects associated at that point in the DAG using the optional interop_handle class.

A library of algorithms based on the C++17 algorithms library was introduced in Section 4.19.4. These algorithms provide a simple way for developers to apply common parallel algorithms using the work-items of a group.

The definition of the sycl::group class was modified to support the new group functions in Section 4.19.5. New member types and variables were added to enable generic programming, and member functions were updated to encapsulate all functionality tied to work-groups in the sycl::group class. See Table 4.85 for details.

The barrier and mem_fence member functions of the nd_item class have been removed. The barrier member function has been replaced by the group_barrier() function, which can be used to synchronize either work-groups or sub-groups. The mem_fence member function has been replaced by the atomic_fence function, which is more closely aligned with std::atomic_thread_fence and offers control over memory ordering and scope.

Changes in the SYCL vec class described in Section 4.16.2:

- operator[] was added;
- unary operator+() and operator-() were added;
- get_count() and get_size() were made static constexpr.

Buffer and local accessors now meet the C++ requirement of ContiguousContainer. This includes (but is not limited to) returning begin and end iterators, specifying a default constructible accessor that can be passed to a

kernel but not dereferenced, and making them equality comparable within kernel functions. accessor::get_size () has been removed to prevent confusion with accessor::size(), and replaced with accessor::byte_size(). All buffer and local accessor size and iterator queries have been marked noexcept.

The device selection now relies on a simpler API based on ranking functions used as device selectors described in Section 4.6.1.1.

A new reduction library consisting of the reduction function, reducer class and parallel_reduce was introduced to simplify the expression of variables with reduction semantics in SYCL kernels. See Section 4.10.2.

Global and constant accessors can now be constructed as placeholders without specifying the access:: placeholder template parameter (which has been deprecated). It is allowed to call handler::require on both placeholder and non-placeholder global and constant buffer accessors, and it is allowed to call it multiple times. accessor::is_placeholder is not constexpr anymore.

The image class has been replaced with unsampled_image and sampled_image classes. The sampler class has been modified to support these changes.

Specified the constness semantics of accessors. const dataT and access::mode::read are semantically equivalent, and having at least one of those parameters part of the accessor type makes the accessor read-only. Defined implicit conversions based on these semantics. Specified default access mode to be access::mode::read for const dataT and access::mode::read_write otherwise. Specified default accessor dimensions to be 1. All of these rules enable most buffer accessor code to only need to use accessor<T> for mutable data and accessor< const T> for const data.

The atomic class from SYCL 1.2.1 and accessors using access::mode::atomic were deprecated in favor of a new atomic_ref interface.

The SYCL exception class hierarchy has been condensed into a single exception type: exception. exception now derives from std::exception. The variety of errors are now provided via error codes, which aligns with the C++ error code mechanism.

The new error code mechanism now also generalizes the previous get_cl_code interface to provide a generic interface way for querying backend-specific error codes.

Default asynchronous error handling behavior is now defined, so that asynchronous errors will cause abnormal program termination even if a user-defined asynchronous handler function is not defined. This prevents asynchronous errors from being silently lost during early stages of application development.

Kernel invocation functions, such as parallel_for, now take kernel functions by const reference. Kernel functions must now have a const-qualified operator(), and are allowed to be copied zero or more times by an implementation. These clarifications allow implementations to have flexibility for specific devices, and define what users should expect with kernel functors. Specifically, kernel functors can not be marked as mutable, and sharing of data between work-items should not be attempted through state stored within a kernel functor.

A new concept called device aspects has been added, which tells the set of optional features a device supports. This new mechanism replaces the has_extension() function and some uses of get_info().

There is a new Chapter 6 which describes how extensions to the SYCL language can be added by vendors and by The Khronos Group.

A queue constructor has been added that takes both a device and context, to simplify interfacing with libraries.

APPENDIX E. WHAT HAS CHANGED FROM PREVIOUS VERSIONS

The parallel_for interface has been simplified in some forms to accept a braced initializer list in place of a range, and to always take item arguments. Kernel invocation functions have also been modified to accept generic lambda expressions. Some implicit conversions have been defined to allow one-dimensional item to convert to scalar types. All of these modifications lead to simpler SYCL code in common use cases.

Some device-specific queries have been renamed to more clearly be "device-specific kernel" get_info queries (info::kernel_device_specific) instead of "work-group" (get_workgroup_info) and sub-group (get_sub_group_info) queries.

A new math array type marray has been defined to begin disambiguation of the multiple possible interpretations of how sycl::vec should be interpreted and implemented.

Changes in SYCL address spaces:

- the address space meaning has been significantly improved;
- the generic address space was introduced;
- behavior of unannotated pointer/reference (raw pointer/reference) is now dependent on the compilation mode. The compiler can either interpret unannotated pointer/reference has addressing the generic address space or to be deduced;
- some ambiguities in the address space deduction were clarified. Notably that deduced type does not affect the user-provided type.

Changes in multi_ptr interface:

- addition of access::address_space::generic_space to represent the generic address space;
- an extra template parameter to allow to select a flavor of the multi_ptr interface. There are now 3 different interfaces:
 - interface exposing undecorated types. Returned pointer and reference are not annotated by an address space;
 - interface exposing decorated types. Returned pointer and reference are annotated by an address space;
 - legacy 1.2.1 interface (deprecated).
- deprecation of the 1.2.1 interface;
- global_ptr, local_ptr, constant_ptr and private_ptr alias take the new extra parameter;
- addition of the address_space_cast free function to cast undecorated pointer to multi_pointer;
- addition of construction/conversion operator for the generic address space;
- removal of the constructor and assignment operator taking an unannotated pointer;
- implicit conversion to a pointer is now deprecated. get should be used instead;
- the return type of the member function get now depends on the selected interface.
- addition of the member function get_raw which returns the underlying pointer as an unannotated pointer;

• addition of the member function get_decorated which returns the underlying pointer as an annotated pointer.

The accessor member function get_pointer now returns a raw pointer. The get_multi_ptr member function was introduced to accessor classes which return the multi_ptr class to the appropriate space.

The cl::sycl::byte has been deprecated and now the C++17 std::byte should be used instead.

References

- [1] Khronos OpenCL Working Group, *The OpenCL Specification, version 1.2.19*, 2012. [Online]. Available: https://www.khronos.org/registry/OpenCL/specs/opencl-1.2.pdf
- [2] International Organization for Standardization (ISO), "Programming Languages C++," Tech. Rep. ISO/IEC 14882:2017, 2017.
- [3] —, Working Draft, Standard for Programming Language C++, 2020. [Online]. Available: http://www.open-std.org/JTC1/SC22/WG21/docs/papers/2020/n4849.pdf
- [4] Khronos OpenCL Working Group, *The OpenCL Specification, Version 2.0*, 2015. [Online]. Available: https://www.khronos.org/registry/OpenCL/specs/opencl-2.0.pdf
- [5] —, *The OpenCL Extension Specification, version 1.2.22, 2012.* [Online]. Available: http://www.khronos.org/registry/cl/specs/opencl-1.2-extensions.pdf

REFERENCES

Glossary

- **accessor** An accessor is a class which allows a SYCL kernel function to access data managed by a buffer or image class. Accessors are used to express the dependencies among the different command groups. For the full description please refer to section [4.7.6]. 30, 31, 37, 38, 109, 110, 124, 142, 144, 263, 265, 266
- SYCL application A SYCL application is a C++ application which uses the SYCL programming model in order to execute kernels on devices. 30, 33, 57, 58, 59, 265, 266, 272, 325, 416, 423, 424, 428
- **application scope** The application scope starts with the construction first SYCL runtime class object and finishes with the destruction of the last one. Application refers to the C++ SYCL application and not the SYCL runtime. 30, 33, 159
- **aspect** A characteristic of a device which determines whether it supports some optional feature. Aspects are always boolean, so a device either has or does not have an aspect. 70, 78, 88, 91, 387, 417
- **async_handler** An asynchronous error handler object is a function class instance providing necessary code for handling all the asynchronous errors triggered from the execution of command groups on a queue, within a context or an associated event. For the full description please refer to section [4.15.2]. 7, 94, 100, 104, 106, 107, 285, 286
- asynchronous error A SYCL asynchronous error is an error occurring after the host API call invoking the error causing action has returned, such that the error cannot be thrown as a typical C++ exception from a host API call. Such errors are typically generated from device kernel invocations which are executed when SYCL task graph dependencies are satisfied, which occur asynchronously from host code execution. For the full description and associated asynchronous error handling mechanisms, please refer to section [4.15]. 46, 94, 100, 104, 106, 107, 285, 286
- **SYCL backend** An implementation of the SYCL programming model using an heterogeneous programming API. A SYCL backend exposes one or multiple SYCL platforms. For example, the OpenCL backend, via the ICD loader, can expose multiple OpenCL platforms. 3, 19, 29, 31, 32, 33, 34, 37, 38, 45, 46, 47, 48, 49, 51, 55, 56, 57, 58, 61, 64, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79, 88, 92, 93, 99, 101, 103, 105, 106, 107, 108, 109, 110, 124, 178, 266, 268, 269, 270, 271, 278, 287, 289, 292, 324, 370, 375, 379, 381, 382, 388, 407, 408, 409, 411, 416, 424, 429
- **buffer** The buffer class manages data for the SYCL C++ host application and the SYCL device kernels. The buffer class may acquire ownership of some host pointers passed to its constructors according to the constructor kind.

The buffer class, together with the accessor class, is responsible for tracking memory transfers and guaranteeing data consistency among the different kernels. The SYCL runtime manages the memory allocations on both the host and the device within the lifetime of the buffer object. For the full description please refer to section [4.7.2]. 31, 37, 42, 48, 49, 109, 142, 158, 263, 265, 423, 429

command A request to execute work that is submitted to a queue such as the invocation of a SYCL kernel

function, the invocation of a host task or an asynchronous copy. 243, 244, 263, 267, 425, 429

- **command group handler** The command group handler class provides the interface for the commands that can be executed inside the command group scope. It is provided as a scoped object to all of the data access requests within the command group scope. For the full description please refer to section [4.10.4]. 159, 239, 244, 248, 260, 262, 263, 266, 424, 427
- **command group function object** A type which is callable with operator() that takes a reference to a command group handler, that defines a command group which can be submitted by a queue. The function object can be a named type, lambda function or std::function. 31, 32, 34, 46, 47, 48, 100, 243, 244, 286, 424
- **command group scope** The command group scope is the function scope defined by the command group function object. The command group command group handler object lifetime is restricted to the command group scope. For more details please see [4.10.3]. 30, 31, 33, 51, 57, 159, 243, 244, 424
- command group In SYCL, the operations required to process data on a device are represented using a command group function object. Each command group function object is given a unique command group handler object to perform all the necessary work required to correctly process data on a device using a kernel. In this way, the group of commands for transferring and processing data is enqueued as a command group on a device for execution. A command group is submitted atomically to a SYCL queue. 11, 33, 34, 35, 36, 37, 38, 39, 94, 104, 105, 108, 109, 145, 248, 260, 261, 262, 263, 266, 272, 423, 424, 425, 426
- **constant memory** A region of global memory that remains constant during the execution of a kernel. The SYCL runtime allocates and initializes memory objects placed into constant memory. 40, 143, 147
- context A context represents the runtime data structures and state required by a SYCL backend API to interact with a group of devices associated with a platform. The context is defined as the sycl::context class, for further details please see [4.6.3]. 31, 32, 34, 48, 72, 73, 74, 247, 262, 263, 264, 266, 271, 272, 276, 278, 279, 280, 281, 402, 408, 409, 424, 427
- **device** A SYCL device encapsulates a number of heterogeneous devices exposed by a SYCL platform from a given SYCL backend. SYCL devices can execute SYCL kernel functions. 29, 30, 31, 32, 33, 34, 48, 71, 73, 74, 79, 80, 88, 91, 92, 93, 94, 103, 262, 263, 264, 266, 271, 275, 276, 279, 280, 402, 403, 414, 423, 424, 425, 426, 427, 428, 429
- **device function** A device function is any function in a SYCL application that can be run on a device. This includes SYCL kernel functions and, recursively, and functions they call. 376
- **device image selection function** A callable object which takes the begin and end iterators of a module pointing to a sequence of device image and returns an iterator to a chosen device image. 262
- **device image** A SYCL device image represents an implementation defined file format that encapsulates the relevant functions, symbols and meta-data to represent the SYCL kernel function set of a module. 247, 262, 271, 272, 274, 276, 277, 278, 279, 280, 424, 427
- **device selector** A way to select a device used in various places. This is a callable object taking a device reference and returning an integer rank. One of the device with the highest positive value is selected. See Section 4.6.1.1 for more details. 48, 66, 67, 68, 69, 70, 75, 77, 97, 98, 417
- **device compiler** A SYCL device compiler is a compiler that produces OpenCL device binaries from a valid SYCL application. For the full description please refer to section [5]. 30, 42, 50, 234, 278, 375, 401

- event A SYCL object that represents the status of an operation that is being executed by the SYCL runtime. 105, 108
- executable A state which a module can be in, representing SYCL kernel functions as an executable. 274, 276, 425

executable module A module that is of module state executable. 272, 276, 277, 278, 427

SYCL file A SYCL C++ source file that contains SYCL API calls. 428

- generic memory Generic memory is a virtual memory region which can represent global, local and private memory region. 41
- **global memory** Global memory is a memory region accessible to all work-items executing on a device. 40, 143, 147, 187
- **global id** As in OpenCL, a global ID is used to uniquely identify a work-item and is derived from the number of global work-items specified when executing a kernel. A global ID is a one, two or three-dimensional value that starts at 0 per dimension. 36, 223, 224, 426, 429
- group A group of work-items within the index space of a SYCL kernel execution, such as a work-group or sub-group. 42, 324, 353, 416, 427
- **group mem-fence** A group mem-fence guarantees that any memory access before a group barrier must complete before continuing to process any data after the barrier. All work-items in the group execute a release fence prior to synchronizing at the barrier, and all work-items in the group execute an acquire fence afterwards. The scope of these fences can be specified by a memory_scope. 359, 425, 427
- **group barrier** A synchronization function within a group of work-items. All the work-items of a group must execute the barrier construct for any work-item continues execution beyond the barrier. Additionally a group barrier performs a group mem-fence. 42, 46, 324, 357, 358, 425, 429
- host Host is the system that executes the C++ application including the SYCL API. 29, 50, 159, 167, 173, 324, 428, 429
- SYCL host device See SYCL host backend. 324, 399, 401, 426
- **host task command** A type of command that can be used inside a command group in order to schedule a native C++ function. 263
- **host task** A command which invokes a native C++ callable, scheduled conforming to SYCL dependency rules. 263, 264, 265, 266, 416, 424

host pointer A pointer to memory on the host. Cannot be accessed directly from a device. 140

SYCL host backend The SYCL host device is a native C++ implementation of a device. It does not have an native handle. It has full SYCL capabilities and reports them through the SYCL information retrieval interface. The SYCL host device is mandatory for every SYCL implementation and is always available, but may not achieve the same performance as a different backend running on the CPU. Any C++ application debugger can be used for debugging SYCL kernels executing on a SYCL host device. 33, 55, 66, 67, 69, 71, 73, 79, 99, 399, 425

Glossary

- id It is a unique identifier of an item in an index space. It can be one, two or three dimensional index space, since the SYCL kernel execution model is an nd-range. It is one of the index space classes. For the full description please refer to section 4.10.1.3. 217, 222, 223, 224, 230, 234, 426, 427
- image Images in SYCL, like buffers, are abstractions of multidimensional structured arrays. Image can refer to unsampled_image and sampled_image. For the full description please refer to section [4.7.3]. 31, 37, 48, 49, 109, 142, 263, 265, 423
- index space classes The OpenCL kernel execution model defines an nd-range index space. The SYCL runtime class that defines an nd-range is the sycl::nd_range, which takes as input the sizes of global and local work-items, represented using the sycl::range class. The kernel library classes for indexing in the defined nd-range are the following classes:
 - sycl::id: The basic index class representing a id.
 - sycl::item: The index class that contains the global id and local id.
 - sycl::nd_item: The index class that contains the global id, local id and the work-group id.
 - sycl::group : The group class that contains the work-group id and the member functions on a work-group.

253

- **input** A state which a module can be in, representing SYCL kernel functions as a source or intermediate representation. 274, 275, 426
- input module A module that is of module state input. 272, 275, 278, 279, 407, 427, 428
- item An item id is an interface used to retrieve the global id, work-group id and local id. For further details see [4.10.1.4]. 220, 225, 226, 227, 426
- **kernel** A SYCL kernel which can be executed on a device, including the SYCL host device. Is created implicitly when defining a SYCL kernel function (see 4.10) but can also be created manually in order to pre-compile SYCL kernel functions. 34, 50, 75, 110, 248, 252, 260, 267, 269, 270, 281, 403, 409, 423, 427
- **kernel invocation command** A type of command that can be used inside a command group in order to schedule a SYCL kernel function, includes single_task, all variants of parallel_for and parallel_for_workgroup. 262, 263, 271, 272, 376, 409, 416, 429
- **kernel handler** A representation of a SYCL kernel function being invoked that is available to the kernel scope. 247
- **kernel name** A kernel name is a class type that is used to assign a name to the kernel function, used to link the host system with the kernel object output by the device compiler. For details on naming kernels please see [5.2]. 50, 248, 249, 250, 251, 252, 283, 429
- **kernel scope** The function scope of the operator() on a SYCL kernel function. Note that any function or member function called from the kernel is also compiled in kernel scope. The kernel scope allows C++ language extensions as well as restrictions to reflect the capabilities of OpenCL devices. The extensions and restrictions are defined in the SYCL device compiler specification. 30, 33, 247, 416, 426
- SYCL kernel function A type which is callable with operator() that takes a id, item, nd-item or work-group

Glossary

which can be passed to kernel enqueue member functions of the command group handler. A SYCL kernel function defines an entry point to a kernel. The function object can be a named trivially copyable type or lambda function. 29, 30, 31, 32, 43, 56, 57, 101, 102, 103, 244, 247, 248, 249, 250, 251, 252, 253, 256, 259, 262, 263, 271, 272, 274, 275, 276, 277, 278, 279, 280, 281, 283, 376, 397, 407, 409, 411, 416, 423, 424, 425, 426, 427, 428, 429

- **local memory** Local memory is a memory region associated with a work-group and accessible only by workitems in that work-group. 41, 142, 143, 167, 170, 171, 257
- local id A unique identifier of a work-item among other work-items of a work-group. 36, 223, 230, 426, 429
- **mem-fence** A memory fence guarantees that any memory access must complete before continuing to process any data after the fence. The sycl::atomic_fence function acts as a fence across all work-items and devices specified by a memory_scope argument, and a group mem-fence acts as a fence across all work-items in a specific group. 42, 324
- **module** A SYCL module represents a set of SYCL kernel functions which can be executed on a number of devices associated with a context. 34, 247, 262, 263, 271, 272, 274, 275, 276, 277, 278, 279, 280, 281, 403, 407, 408, 409, 416, 424, 425, 426, 427
- **module state** A SYCL module state represents an the state abstract state of a module and therefore it's capabilities in the SYCL programming API. Can be either a input module, object module or an executable module. 274, 275, 276, 278, 281, 425, 426, 427
- **native-specialization constant** A specialization constant which is natively supported by the device image which contains it. 272, 277
- native backend object An opaque object defined by a specific backend that represents a high-level SYCL object on said backend. There is no guarantee of having native backend objects for all SYCL types. 48, 57, 58, 59, 60, 61, 68, 75, 94, 194, 263, 264, 265, 266
- **nd-item** A unique identifier representing a single work-item and work-group within the index space of a SYCL kernel execution. Can be one, two or three dimensional. In the SYCL interface a nd-item is represented by the nd_item class (see Section 4.10.1.5). 103, 250, 252, 426, 427
- nd-range A representation of the index space of a SYCL kernel execution, the distribution of work-items within into work-groups. Contains a range specifying the number of global work-items, a range specifying the number of local work-items and a id specifying the global offset. Can be one, two or three dimensional. The minimum size of each range within the nd-range is 1 per dimension. In the SYCL interface an nd-range is represented by the nd_range class (see Section 4.10.1.2). 36, 103, 223, 230, 248, 250, 252, 256, 257, 426, 427
- **object** A state which a module can be in, representing SYCL kernel functions as a non-executable object. 274, 275, 427

object module A module that is of module state object. 272, 275, 427

platform The host together or a collection of devices managed by the OpenCL framework that allow an application to share resources and execute kernels on devices in the platform. A SYCL application can target one

Glossary

or multiple OpenCL platforms provided by OpenCL device vendors [1]. 32, 33, 34, 48, 71, 74, 402, 423, 424

- **private memory** A region of memory private to a work-item. Variables defined in one work-item's private memory are not visible to another work-item. [1]. The sycl::private_memory class provides access to the work-item's private memory for the hierarchical API as it is described at [4.10.7.3]. 41, 258
- **queue** A SYCL command queue is an object that holds command groups to be executed on a SYCL device. SYCL provides a heterogeneous platform integration using device queue, which is the minimum requirement for a SYCL application to run on a SYCL device. For the full description please refer to section [4.6.5]. 31, 32, 33, 34, 47, 48, 103, 262, 263, 264, 266, 402, 403, 416, 423, 424
- **range** A representation of a number of work-items or work-group within the index space of a SYCL kernel execution. Can be one, two or three dimensional. In the SYCL interface a work-group is represented by the group class (see Section 4.10.1.7). 427
- **reduction** An operation that produces a single value by combining multiple values in an unspecified order using a binary operator. If the operator is non-associative or non-commutative, the behavior of a reduction may be non-deterministic. 44, 234, 235, 237, 417
- rule of zero For a given class, if the copy constructor, move constructor, copy assignment operator, move assignment operator and destructor would all be inlined, public and defaulted, none of them should be explicitly declared . 62, 63
- rule of five For a given class, if at least one of the copy constructor, move constructor, copy assignment operator, move assignment operator or destructor is explicitly declared, all of them should be explicitly declared . 62, 63
- SYCL runtime A SYCL runtime is an implementation of the SYCL API specification. The SYCL runtime manages the different OpenCL platforms, devices, contexts as well as memory handling of data between host and OpenCL contexts to enable semantically correct execution of SYCL programs. 32, 39, 45, 46, 47, 48, 49, 50, 55, 56, 57, 58, 59, 60, 62, 64, 66, 94, 104, 109, 111, 118, 121, 123, 125, 127, 129, 130, 131, 132, 135, 136, 138, 139, 140, 142, 148, 158, 159, 167, 173, 230, 244, 247, 249, 250, 251, 252, 253, 263, 264, 265, 266, 267, 272, 276, 284, 285, 291, 326, 375, 377, 401, 403, 409, 411, 413, 416, 423, 424, 426
- **SMCP** The single-source multiple compiler-passes (SMCP) technique allows a single source file to be parsed by multiple compilers for building native programs per compilation target. For example, a standard C++ CPU compiler for targeting host will parse the SYCL file to create the C++ SYCL application which offloads parts of the computation to other devices. A SYCL device compiler will parse the same source file and target only SYCL kernels. 27, 28, 50, 92, 375, 387
- **specialization id** An identifier which represents a reference to a specialization constant both in the SYCL application for setting the value prior to the compilation of an input module and in a SYCL kernel function for retrieving the value during invocation. 272, 275
- **specialization constant** A constant variable where the value is not known until compilation of the SYCL kernel function. 247, 263, 272, 275, 277, 278, 283, 416, 427, 428
- string kernel name The name of a SYCL kernel function in string form, this can be the name of a kernel function created via interop or a string form of a type kernel name. 277, 280

sub-group barrier A group barrier for all work-items in a sub-group. 46

- sub-group The SYCL sub-group (sycl::sub_group class) is a representation of a collection of related workitems within a work-group that execute concurrently, and which may make independent forward progress with respect to other sub-groups in the same work-group. For further details for the sycl::sub_group class see [4.10.1.8]. 42, 223, 233, 234, 416, 425, 429
- SYCL C++ template library The template library is a set of C++ templated classes which provide the programming interface to the SYCL developer. 401
- SYCL backend API The exposed API for writing SYCL code against a given SYCL backend. 3, 29, 34, 35, 48, 50, 57
- **type kernel name** The name of a SYCL kernel function in type form, this can be either a kernel name provided to akernel invocation command or the type of a function object use as a SYCL kernel function. 280, 428
- **USM** Unified Shared Memory (USM) provides a pointer-based alternative to the buffer programming model. USM enables:
 - easier integration into existing code bases by representing allocations as pointers rather than buffers, with full support for pointer arithmetic into allocations
 - fine-grain control over ownership and accessibility of allocations, to optimally choose between performance and programmer convenience;
 - a simpler programming model, by automatically migrating some allocations between SYCL devices and the host.

See Section 4.8. 27, 42, 196, 197, 377

work-group barrier A group barrier for all work-items in a work-group. 46, 255, 256, 258, 324

- **work-group id** As in OpenCL, SYCL kernels execute in work-groups. The group ID is the ID of the work-group that a work-item is executing within. A group ID is an one, two or three dimensional value that starts at 0 per dimension. 36, 222, 230, 426, 429
- work-group The SYCL work-group (sycl::group class) is a representation of a collection of related work-items that execute on a single compute unit. The work-items in the group execute the same kernel-instance and share local memory and work-group functions [1]. For further details for the sycl::group class see [4.10.1.7]. 36, 42, 167, 223, 224, 228, 230, 231, 234, 248, 256, 257, 416, 425, 426, 427, 428, 429
- work-item The SYCL work-item is a representation of a work-item among a collection of parallel executions of a kernel invoked on a device by a command. A work-item is executed by one or more processing elements as part of a work-group executing on a compute unit. A work-item is distinguished from other work-items by its global id or the combination of its work-group id and its local id within a work-group [1]. 36, 42, 167, 221, 222, 224, 225, 256, 257, 258, 416, 425, 427, 428, 429