



SYCL[™] Specification

SYCL integrates OpenCL devices with modern C++

Version 1.2

Revision Date: 2015-05-08

Khronos OpenCL Working Group — SYCL subgroup

Editors: Lee Howes and Maria Rovatsou

Contents

1	Introduction	11
2	SYCL Architecture	13
2.1	Overview	13
2.2	The SYCL Platform Model	13
2.2.1	Platform Mixed Version Support	14
2.3	SYCL Execution Model	14
2.3.1	Execution Model: Queues, Command Groups and Contexts	15
2.4	Memory Model	16
2.4.1	Access to memory	17
2.4.2	Memory consistency	18
2.4.3	Atomic operations	18
2.5	The SYCL programming model	18
2.5.1	Basic data parallel kernels	19
2.5.2	Work-group data parallel kernels	19
2.5.3	Hierarchical data parallel kernels	19
2.5.4	Kernels that are not launched over parallel instances	20
2.5.5	Synchronization	20
2.5.6	Error handling	21
2.5.7	Scheduling of kernels and data movement	21
2.5.8	Managing object lifetimes	23
2.5.9	Device discovery and selection	24
2.5.10	Interfacing with OpenCL	24
2.6	Anatomy of a SYCL application	26
2.7	Memory objects	27
2.8	SYCL for OpenCL Framework	28
2.9	SYCL device compiler	29
2.9.1	Building a SYCL program	29
2.9.2	Naming of kernels	29
2.10	Language restrictions in kernels	30
2.10.1	SYCL Linker	30
2.10.2	Functions and datatypes available in kernels	31
2.11	Execution of kernels on the SYCL host device	31
2.12	Example SYCL application	31
3	SYCL Programming Interface	34
3.1	Header files and namespaces	34
3.2	C++ Standard library classes required for the interface	34
3.3	SYCL runtime classes	35
3.3.1	Device selection class	35
3.3.2	Platform class	37
3.3.2.1	Platform information descriptors	39
3.3.3	Context class	39

3.3.3.1	Context interface	40
3.3.3.2	Context information descriptors	42
3.3.4	Device class	42
3.3.4.1	Device interface	43
3.3.4.2	Device information descriptors	45
3.3.5	Queue class	55
3.3.5.1	Queue interface	55
3.3.5.2	Queue information descriptors	59
3.3.5.3	Queue error handling	60
3.3.6	Event class for OpenCL interoperability	60
3.3.6.1	Event information and profiling descriptors	62
3.4	Data access and storage in SYCL	64
3.4.1	Host allocation	64
3.4.1.1	Default Allocators	64
3.4.1.2	Map Allocator	65
3.4.2	Buffers	65
3.4.2.1	Buffer Interface	66
3.4.2.2	Buffer Synchronization Rules	69
3.4.3	Images	71
3.4.3.1	Image Interface	71
3.4.3.2	Image Synchronization Rules	75
3.4.4	Sharing Host Memory With The SYCL Data Management Classes	75
3.4.4.1	Default behavior	76
3.4.4.2	SYCL ownership of the host memory	76
3.4.4.3	Shared SYCL ownership of the host memory	77
3.4.5	Synchronisation Primitives	77
3.4.6	Accessors	78
3.4.6.1	Access modes	78
3.4.6.2	Access targets	79
3.4.6.3	Accessor class	79
3.4.6.4	Buffer accessors	84
3.4.6.5	Image accessors	85
3.4.6.6	Local accessors	85
3.4.6.7	Host accessors	85
3.4.6.8	Accessor capabilities and restrictions	86
3.4.7	Address space classes	88
3.4.7.1	Explicit pointer classes	88
3.4.7.2	Multi-pointer class	97
3.4.8	Samplers	102
3.5	Expressing parallelism through kernels	103
3.5.1	Ranges and index space identifiers	103
3.5.1.1	range class	103
3.5.1.2	nd_range class	108
3.5.1.3	id class	109
3.5.1.4	item class	114
3.5.1.5	nd_item class	115
3.5.1.6	group class	117
3.5.2	Command group scope	118
3.5.3	SYCL functions for invoking kernels	119
3.5.3.1	single_task invoke	119
3.5.3.2	parallel_for invoke	120

3.5.3.3	Parallel For hierarchical invoke	122
3.5.3.4	Command group handler class	124
3.5.4	Kernel class	128
3.5.5	Program class	129
3.5.6	Defining kernels	132
3.5.6.1	Defining kernels as functors	132
3.5.6.2	Defining kernels as lambda functions	133
3.5.6.3	Defining kernels using program objects	134
3.5.6.4	Defining kernels using OpenCL C kernel objects	135
3.5.7	Rules for parameter passing to kernels	135
3.6	Error handling	136
3.6.1	Error Handling Rules	136
3.6.2	Exception Class Interface	137
3.7	Data types	140
3.7.1	Scalar data types	140
3.7.2	Vector types	142
3.8	Synchronization and atomics	148
3.9	Stream class	154
3.10	SYCL built-in functions for SYCL host and device	155
3.10.1	Description of the built-in types available for SYCL host and device	155
3.10.2	Work-item functions	157
3.10.3	Math functions	157
3.10.4	Integer functions	162
3.10.5	Common functions	163
3.10.6	Geometric Functions	164
3.10.7	Relational functions	166
3.10.8	Vector data and store functions	167
3.10.9	Synchronization Functions	168
3.10.10	printf function	168
4	SYCL Support of Non-Core OpenCL Features	169
4.1	Enable extensions in a SYCL kernel	169
4.2	Half Precision Floating-Point	170
4.3	Writing to 3D image memory objects	171
4.4	Interoperability with OpenGL	171
4.4.1	OpenCL/OpenGL extensions to the context class	171
4.4.2	Sharing SYCL/OpenGL memory objects	172
4.4.2.1	SYCL/OpenGL extensions to SYCL buffer	173
4.4.2.2	SYCL/OpenGL extensions to SYCL image	173
4.4.2.3	SYCL/OpenGL extensions to SYCL accessors	175
4.4.2.4	SYCL/OpenGL extensions to SYCL events	176
4.4.2.5	Extension for depth and depth-stencil images	176
5	SYCL Device Compiler	177
5.1	Offline compilation of SYCL source files	177
5.2	Naming of kernels	177
5.3	Language restrictions for kernels	178
5.4	Compilation of functions	179
5.5	Built-in scalar data types	179
5.6	Preprocessor directives and macros	180
5.7	Attributes	181

5.8	Address-space deduction	181
5.9	SYCL offline linking	182
5.9.1	SYCL functions and methods linkage	182
5.9.2	Offline linking with OpenCL C libraries	182
A	Glossary	183
B	Interface of SYCL Classes in Full	186
B.1	Explicit pointer classes	186
B.2	Multi pointer pointer class	191
B.3	range class	192
B.4	id class	195
B.5	vec class	196
C	Interface of Memory Object Information Descriptors	199
C.1	Platform Information Descriptors	199
C.2	Context Information Descriptors	199
C.3	Device Information Descriptors	200
C.4	Queue Information Descriptors	202
C.5	Kernel Information Descriptors	203
C.6	Program Information Descriptors	203
C.7	Event Information Descriptors	204
	References	205

List of Tables

3.1	Constructors of the <code>device_selector</code> class	36
3.2	Methods for the <code>device_selector</code> class	36
3.3	Standard device selectors included with all SYCL implementations.	36
3.3	Standard device selectors included with all SYCL implementations.	37
3.4	Constructors of <code>platform</code> class	38
3.5	Methods of <code>platform</code> class	39
3.6	Platform information descriptors.	39
3.7	Constructors of the <code>context</code> class.	41
3.8	Methods of context class	42
3.9	Context information descriptors	42
3.10	Constructors of the <code>device</code> class	44
3.11	Methods of the <code>device</code> class	44
3.11	Methods of the <code>device</code> class	45
3.12	Device information descriptors.	45
3.12	Device information descriptors.	46
3.12	Device information descriptors.	47
3.12	Device information descriptors.	48
3.12	Device information descriptors.	49
3.12	Device information descriptors.	50
3.12	Device information descriptors.	51
3.12	Device information descriptors.	52
3.12	Device information descriptors.	53
3.12	Device information descriptors.	54
3.12	Device information descriptors.	55
3.13	Constructors of the <code>queue</code> class.	57
3.13	Constructors of the <code>queue</code> class.	58
3.14	Methods for class <code>queue</code>	58
3.14	Methods for class <code>queue</code>	59
3.15	Queue information descriptors	59
3.15	Queue information descriptors	60
3.16	Constructors for the <code>event</code> class	61
3.17	Methods for the <code>event</code> class	61
3.17	Methods for the <code>event</code> class	62
3.18	Event class information descriptors.	63
3.19	Event class profiling information descriptors.	63
3.19	Event class profiling information descriptors.	64
3.20	SYCL Default Allocators	65
3.21	Constructors for the <code>buffer</code> class.	67
3.21	Constructors for the <code>buffer</code> class.	68
3.22	Methods for the <code>buffer</code> class.	68
3.22	Methods for the <code>buffer</code> class.	69
3.23	Constructors for the <code>image</code> class.	73
3.23	Constructors for the <code>image</code> class.	74

3.24	Methods of the <code>image</code> class.	74
3.24	Methods of the <code>image</code> class.	75
3.25	Enumeration of access modes available to accessors.	79
3.26	Enumeration of access modes available to accessors.	79
3.27	Accessor constructors.	82
3.28	Methods for the <code>accessor</code> class.	83
3.28	Methods for the <code>accessor</code> class.	84
3.29	Description of all the accessor types and modes with their valid combinations for buffers and local memory	86
3.30	Description of all the accessor types and modes with their valid combinations for images	86
3.31	Description of the accessor to accessor conversions allowed	87
3.32	Description of the pointer classes	88
3.33	Constructors for <code>global_ptr</code> explicit pointer class.	90
3.34	Operators on the <code>global_ptr</code> explicit pointer class.	90
3.35	Constructors for <code>constant_ptr</code> explicit pointer class.	91
3.36	Operators on the <code>constant_ptr</code> explicit pointer class.	91
3.37	Constructors for <code>local_ptr</code> explicit pointer class.	91
3.38	Operators on the <code>local_ptr</code> explicit pointer class.	91
3.39	Constructors for <code>private_ptr</code> explicit pointer class.	92
3.40	Operators on the <code>private_ptr</code> explicit pointer class.	92
3.41	Non-member functions of the explicit pointer classes.	92
3.41	Non-member functions of the explicit pointer classes.	93
3.41	Non-member functions of the explicit pointer classes.	94
3.41	Non-member functions of the explicit pointer classes.	95
3.41	Non-member functions of the explicit pointer classes.	96
3.41	Non-member functions of the explicit pointer classes.	97
3.42	Constructors for <code>multi_ptr</code> class	99
3.43	Methods of <code>multi_ptr</code> class	99
3.43	Methods of <code>multi_ptr</code> class	100
3.44	Non-member functions of the <code>multi_ptr</code> class.	100
3.44	Non-member functions of the <code>multi_ptr</code> class.	101
3.45	Constructors for the <code>sampler</code> class.	102
3.45	Constructors for the <code>sampler</code> class.	103
3.46	Methods for the <code>sampler</code> class.	103
3.47	Constructors for the <code>range</code> class.	104
3.48	Methods for the <code>range</code> class.	104
3.48	Methods for the <code>range</code> class.	105
3.49	Non-member functions for the <code>range</code> class.	105
3.49	Non-member functions for the <code>range</code> class.	106
3.49	Non-member functions for the <code>range</code> class.	107
3.49	Non-member functions for the <code>range</code> class.	108
3.50	Constructors for the <code>nd_range</code> class.	109
3.51	Methods for the <code>nd_range</code> class.	109
3.52	Constructors for the <code>id</code> class.	110
3.53	Methods for the <code>id</code> class.	110
3.53	Methods for the <code>id</code> class.	111
3.54	Non-member functions for the <code>id</code> class.	111
3.54	Non-member functions for the <code>id</code> class.	112
3.54	Non-member functions for the <code>id</code> class.	113
3.54	Non-member functions for the <code>id</code> class.	114
3.55	Methods for the <code>item</code> class.	115

3.56	Methods for the <code>nd_item</code> class.	116
3.56	Methods for the <code>nd_item</code> class.	117
3.57	Methods for the <code>group</code> class.	118
3.58	Constructors for the command group <code>handler</code> class	125
3.59	Methods for the command group <code>handler</code> class	126
3.59	Methods for the command group <code>handler</code> class	127
3.60	<code>kernel</code> class constructors	129
3.61	Methods for the <code>kernel</code> class.	129
3.62	Kernel class information descriptors.	129
3.63	Constructors for the <code>program</code> class	131
3.64	Methods for the <code>program</code> class	131
3.64	Methods for the <code>program</code> class	132
3.65	Program class information descriptors.	132
3.66	Methods of the <code>exception</code> class.	138
3.67	Methods of the <code>cl_exception</code> class.	139
3.68	Methods of the <code>exception_list</code>	139
3.69	Exceptions types that derive from the <code>cl::sycl::runtime_error</code> class	139
3.70	Exception types that derive from the <code>cl::sycl::device_error</code> class	139
3.70	Exception types that derive from the <code>cl::sycl::device_error</code> class	140
3.71	SYCL compiler fundamental integral datatypes	140
3.71	SYCL compiler fundamental integral datatypes	141
3.72	SYCL compiler fundamental floating point datatypes	141
3.73	SYCL compiler OpenCL interoperability scalar datatypes	141
3.74	Generic type name description for <code>genVector</code> , which serves as a description for all valid types of OpenCL/SYCL interoperability vectors.	143
3.75	Constructors for the <code>vec</code> class	144
3.76	Methods for the <code>vec</code> class	144
3.76	Methods for the <code>vec</code> class	145
3.76	Methods for the <code>vec</code> class	146
3.76	Methods for the <code>vec</code> class	147
3.77	Methods available on an object of type <code>atomic<T></code>	151
3.77	Methods available on an object of type <code>atomic<T></code>	152
3.78	Global functions available on atomic types.	153
3.78	Global functions available on atomic types.	154
3.80	Generic type name description, which serves as a description for all valid types of parameters to kernel functions. [1]	155
3.80	Generic type name description, which serves as a description for all valid types of parameters to kernel functions. [1]	156
3.81	Math functions which work on SYCL Host and device. They correspond to Table 6.7 of the OpenCL 1.2 specification [1]	157
3.81	Math functions which work on SYCL Host and device. They correspond to Table 6.7 of the OpenCL 1.2 specification [1]	158
3.81	Math functions which work on SYCL Host and device. They correspond to Table 6.7 of the OpenCL 1.2 specification [1]	159
3.81	Math functions which work on SYCL Host and device. They correspond to Table 6.7 of the OpenCL 1.2 specification [1]	160
3.82	Native functions which work on SYCL Host and device, are available in the <code>cl::sycl::native</code> namespace. They correspond to Table 6.9 of the OpenCL 1.2 specification [1]	160
3.82	Native functions which work on SYCL Host and device, are available in the <code>cl::sycl::native</code> namespace. They correspond to Table 6.9 of the OpenCL 1.2 specification [1]	161

3.83	Integer functions which work on SYCL Host and device, are available in the <code>cl::sycl</code> namespace. They correspond to Table 6.10 of the OpenCL 1.2 specification [1]	162
3.83	Integer functions which work on SYCL Host and device, are available in the <code>cl::sycl</code> namespace. They correspond to Table 6.10 of the OpenCL 1.2 specification [1]	163
3.84	Common functions which work on SYCL Host and device, are available in the <code>cl::sycl</code> namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1]	163
3.84	Common functions which work on SYCL Host and device, are available in the <code>cl::sycl</code> namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1]	164
3.85	Geometric functions which work on SYCL Host and device, are available in the <code>cl::sycl</code> namespace. They correspond to Table 6.13 of the OpenCL 1.2 specification [1]	164
3.85	Geometric functions which work on SYCL Host and device, are available in the <code>cl::sycl</code> namespace. They correspond to Table 6.13 of the OpenCL 1.2 specification [1]	165
3.86	Relational functions which work on SYCL Host and device, are available in the <code>cl::sycl</code> namespace. They correspond to Table 6.14 of the OpenCL 1.2 specification [1]	166
3.86	Relational functions which work on SYCL Host and device, are available in the <code>cl::sycl</code> namespace. They correspond to Table 6.14 of the OpenCL 1.2 specification [1]	167
4.1	SYCL support for OpenCL 1.2 API extensions.	169
4.1	SYCL support for OpenCL 1.2 API extensions.	170
4.2	Generic type name description for all valid types of kernel function parameters. [1]	170
4.3	Extended elementary functions which work on SYCL host and device.	171
4.4	Additional optional properties for creating context for SYCL/OpenGL sharing.	172
4.5	Extended <code>context</code> class interface.	172
4.6	Extended constructors for the <code>buffer</code> class.	173
4.7	Extended <code>buffer</code> class interface	173
4.8	Additional optional <code>image</code> class constructors.	173
4.8	Additional optional <code>image</code> class constructors.	174
4.9	Additional optional <code>image</code> class method.	174
4.10	Mapping of GL internal format to CL image format (reference: [2, table 9.4])	175
4.11	Enumerator description for <code>access::target</code>	175
4.12	Additional optional class constructors for <code>event</code> class.	176
4.13	Additional optional class method for <code>event</code> class.	176
5.1	SYCL compiler fundamental scalar datatypes	180

Copyright (c) 2013-2015 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, SYCL, SPIR, WebGL, EGL, COLLADA, StreamInput, OpenVX, OpenKCam, glTF, OpenKODE, OpenVG, OpenWF, OpenSL ES, OpenMAX, OpenMAX AL, OpenMAX IL and OpenMAX DL are trademarks and WebCL is a certification mark of the Khronos Group Inc. OpenCL is a trademark of Apple Inc. and OpenGL and OpenML are registered trademarks and the OpenGL ES and OpenGL SC 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.

Acknowledgements

Editors: Lee Howes, Qualcomm; Maria Rovatsou, Codeplay

Contributors:

- Eric Berdahl, Adobe
- Shivani Gupta, Adobe
- David Neto, Altera
- Ronan Keryell, AMD
- Brian Sumner, AMD
- Balázs Keszthelyi, Broadcom
- Gordon Brown, Codeplay
- Paul Keir, Codeplay
- Ralph Potter, Codeplay
- Ruyman Reyes, Codeplay
- Andrew Richards, Codeplay
- Maria Rovatsou, Codeplay
- Allen Hux, Intel
- Matt Newport, EA
- Lee Howes, Qualcomm
- Chu-Cheow Lim, Qualcomm
- Jack Liu, Qualcomm

1. Introduction

SYCL is a C++ programming model for OpenCL. SYCL builds on the underlying concepts, portability and efficiency of OpenCL while adding much of the ease of use and flexibility of C++. Developers using SYCL are able to write standard C++ code, with many of the techniques they are accustomed to, such as inheritance and templating. At the same time developers have access to the full range of capabilities of OpenCL both through the features of the SYCL libraries and, where necessary, through interoperation with code written directly to the OpenCL APIs.

SYCL implements a shared source design which offers the power of source integration while allowing toolchains to remain flexible. The shared source design supports embedding of code intended to be compiled for an OpenCL device, for example a GPU, inline with host code. This embedding of code offers three primary benefits:

Simplicity For novice programmers, 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. 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 OpenCL 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 shared source 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.

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 an OpenCL 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 optimised tool-chains.

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 OpenCL, a standard C++ compiler can compile the SYCL programs and

they will run correctly on 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 retains the execution model, runtime feature set and device capabilities of the underlying OpenCL standard. The OpenCL C specification 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.

The use of C++ features such as templates and inheritance on top of the OpenCL 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 libraries that build simple, yet efficient, interfaces to offer more developers access to OpenCL capabilities and devices. SYCL is intended to serve as a foundation for innovation in programming models for heterogeneous systems, that builds on an open and widely implemented standard foundation in the form of OpenCL.

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

- The hierarchical parallelism syntax offers a way of expressing the data-parallel OpenCL execution model in an easy-to-understand 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.
- 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 are already part of the OpenCL execution model to be built up easily and safely by SYCL programmers.

To summarize, SYCL enables OpenCL kernels to be written inside C++ source files. This means that software developers can develop and use generic algorithms and data structures using standard C++ template techniques, while still supporting the multi-platform, multi-device heterogeneous execution of OpenCL. 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 heterogeneous processing innovation.

2. SYCL Architecture

This chapter builds on the structure of the OpenCL specification's architecture chapter to explain how SYCL overlays the OpenCL specification and inherits its capabilities and restrictions as well as the additional features it provides on top of OpenCL 1.2.

2.1 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 OpenCL device or on the host CPU.

The terminology used for SYCL inherits that of OpenCL with some SYCL-specific additions. The code that can run on either an OpenCL device or host CPU is called a *kernel*. To ensure maximum backward-compatibility, a software developer can produce a program that mixes standard OpenCL C kernels and OpenCL API code with SYCL code and expect fully compatible interoperation.

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 OpenCL standard. 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 applications developers who want to separate low-level highly-tuned algorithms or data structures that work on heterogeneous systems from higher-level software development. OpenCL developers can produce templated algorithms that are easily usable by developers in other fields.

2.2 The SYCL Platform Model

The SYCL platform model is based on the OpenCL platform model, but there are a few additional abstractions available to programmers.

The model consists of a host connected to one or more OpenCL devices. An OpenCL 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. A SYCL application runs on a host according to the standard C++ CPU execution model. The SYCL application submits **command groups functors** to **queues**, which execute either on OpenCL devices or on the SYCL host device.

When a SYCL implementation executes kernels on an OpenCL device, it achieves this by enqueueing 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 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.

When a SYCL implementation executes kernels on the host device, it is free to use whatever parallel execution facilities are available on the host, as long as it executes within the semantics of the kernel execution model defined by OpenCL.

2.2.1 Platform Mixed Version Support

OpenCL is designed to support devices with different capabilities under a single platform. This includes devices which conform to different versions of the OpenCL specification and devices which support different extensions to the OpenCL specification. There are three important sets of capabilities to consider for a SYCL device: the platform version, the version of a device and the extensions supported.

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 `cl::sycl::device::get_info()` method. Examples of attributes associated with the device version are resource limits and information about functionality beyond the core OpenCL 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.

In OpenCL, a device has a *language version*. In SYCL, the source language is compiled offline, so the language version is not available at runtime. Instead, the SYCL language version is available as a compile-time macro: `CL_SYCL_LANGUAGE_VERSION`.

2.3 SYCL Execution Model

Execution of a SYCL program occurs in two parts: *kernels* that execute on either the host CPU, or one or more *OpenCL devices*, and a *host program* that executes on the host CPU. The host program defines the context for the kernels and manages their execution. Like OpenCL, SYCL is capable of running kernels on multiple device types. However, SYCL builds on top of OpenCL due to the integration into a host toolchain by providing an ability to run kernel code directly on the CPU without interacting with an OpenCL runtime. This is distinct from running on the CPU via an OpenCL device and can be used when no OpenCL platform is available on the machine.

In OpenCL, *queues* contain *commands*, which can include data transfer operations, synchronization commands, or *kernels* submitted for execution. In SYCL, the commands are grouped together into a functor object called *command group*. Command groups associate sets of data movement operations with kernels that will be enqueued together on an underlying OpenCL queue, through the command group *handler* object which is created for every command group functor. These data transfer operations may be needed to make available data that the kernel needs or to return its results to other devices.

When a kernel is submitted for execution by the host, 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 `nd_range<N>` and `id<N>` are each N -dimensional arrays of integers. The iteration space defined via an `range<N>` is an N -dimensional index space starting at the `nd_range`'s global offset and being of the size of 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 workgroup size is left undefined. A kernel invoked over a `range<N>`, instead of an `nd_range<N>` is executed within an iteration space of undefined workgroup size. In this case, less information is available to each work-item through the simpler `item<N>` class.

2.3.1 Execution Model: Queues, Command Groups and Contexts

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 `cl::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 runtime while data movement between contexts must involve the host. A given context can only wrap devices owned by a single platform. In SYCL, a context resource is accessible through by a `cl::sycl::context` object.

3. *Devices*: Platforms provide one or more devices for executing kernels. In SYCL, a device is accessible through a `cl::sycl::device` object. SYCL provides the abstract `cl::sycl::device_selector` class which the user can subclass to define how the runtime should select the best device from all available platforms for the user to use. For ease of use, SYCL provides a set of predefined concrete `device_selector` instances that select devices based on common criteria, such as type of device. SYCL, unlike OpenCL, defines a host device, which means any work that uses the host device will execute on the host and not on any OpenCL device.
4. *Command groups*: SYCL groups OpenCL *commands* into a functor given a unique command group `handler` object to perform all the necessary work required to correctly process host data on a device using a kernel. In this way, they group the commands of transferring and processing these data in order to enqueue them on a device for execution. Command groups are submitted to a SYCL queue.
5. *Kernels*: The SYCL functions that run on SYCL devices (i.e. either an OpenCL device, or the host) are defined as C++ functors or lambda functions. In SYCL, all kernels must have a *name*, which must be a globally-accessible C++ typename. This is required to enable kernels compiled with one compiler to be linked to host code compiled with a different compiler. For functors, the typename of the functor is sufficient as the kernel *name*, but for C++11 lambda functions, the user must provide a user-defined *name*.
6. *Program Objects*: 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 `cl::sycl::program` class.
7. *Command-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 `cl::sycl::queue` objects.

The command-queue schedules commands submitted to it for execution on a device. Commands launched by the host execute asynchronously with respect to the host thread, and not necessarily ordered with respect to each other. It is the responsibility of the SYCL implementation to ensure that the different commands execute in an order which preserves SYCL semantics. This means that a SYCL implementation must map, move or copy data between host and device memory, execute kernels and perform synchronization between different queues, devices and the host in a way that matches the semantics defined in this specification. If a command group runs on an OpenCL device, then this is expected to be achieved by enqueueing the right memory and synchronization commands to the queue to ensure correct execution. If a command group runs on the host device, then this is expected to be achieved by host-specific synchronization as well as by ensuring that no OpenCL device is simultaneously using any required data.

In OpenCL, queues can operate using in-order execution or out-of-order execution. In SYCL, the implementation must add synchronization commands to queues to provide when possible out-of-order execution ordering, regardless of whether the underlying OpenCL queue is in-order or out-of-order.

2.4 Memory Model

Work-items executing in a kernel have access to four distinct memory regions:

- *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 a distinct memory region shared between 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 that work-group. Work-group-level visibility allows local memory to be implemented as dedicated regions of memory on an OpenCL device 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.

The application running on the host uses SYCL buffer objects using instances of the `cl::sycl::buffer` class to allocate memory in the global address space, or can allocate specialized image memory using the `cl::sycl::image` class. In OpenCL, a memory object is attached to a specific context. In SYCL, a `cl::sycl::buffer` or `cl::sycl::image` object can encapsulate multiple underlying OpenCL memory objects and host memory allocations to enable the same buffer or image to be shared between multiple devices in different contexts, and hence different platforms. It is the responsibility of the SYCL implementation to ensure that a buffer or image object shared between multiple OpenCL contexts is moved between contexts using the correct synchronization and copy commands to preserve SYCL memory ordering semantics.

2.4.1 Access to memory

To access global memory inside a kernel, the user must create a `cl::sycl::accessor` object which parameterizes the type of access the kernel requires. 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, write-only or read-write. 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. Atomic access can also be requested on an accessor which allows `cl::sycl::atomic` classes to be used via the accessor.

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 `cl::sycl::local_accessor` object to the kernel as a parameter, or can define a variable in workgroup scope inside `cl::sycl::parallel_for_work_group`.

Any variable defined inside a `cl::sycl::parallel_for` scope or `cl::sycl::parallel_for_work_item` scope will be allocated in private memory. Variables defined in functions called from workgroup scope (i.e. `cl::sycl::parallel_for_work_group`) will also be local, while variables defined in functions called from workitem scope (i.e. `cl::sycl::parallel_for` or `cl::sycl::parallel_for_work_item`) will be allocated in private memory.

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

Within kernels, accessors can be implicitly cast to C++ pointer types. The pointer types will contain a compile-time deduced address space. So, for example, if an accessor to global memory is cast to a C++ pointer, 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 mechanism.

When developers need to explicitly state the memory 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: `cl::sycl::local_ptr`, `cl::sycl::global_ptr`, `cl::sycl::private_ptr`, or `cl::sycl::constant_ptr`. An accessor declared with one address space can be implicitly cast to an explicit pointer class for the same address space. 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 `cl::sycl::multi_ptr` class is defined which is templated via a compile-time constant enumerator value to specify the address space.

2.4.2 Memory consistency

OpenCL uses a relaxed memory consistency model, i.e. the state of memory visible to a work-item is not guaranteed to be consistent across the collection of work-items at all times. This also applies to SYCL kernels.

As in OpenCL, within a work-item memory has load/store consistency. Both local memory and global memory may be made consistent across work-items in a single work-group through use of a work-group barrier operation with appropriate flags. There are no guarantees of memory consistency between different work-groups executing a kernel or between different kernels during their execution.

Memory consistency for `cl::sycl::buffer` and `cl::sycl::image` objects shared between enqueued commands is enforced at synchronization points derived from completion of enqueued commands. Consistency of such data between the OpenCL runtime and the host program is ensured via copy commands or map and unmap operations.

2.4.3 Atomic operations

Atomic operations can be performed on memory in buffers. The range of atomic operations available on a specific OpenCL device is limited by the atomic capabilities of that device. The `cl::sycl::atomic<T>` must be used for elements of a buffer to provide safe atomic access to the buffer from device code.

2.5 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 C++ features used in SYCL are a subset of the C++14 standard features. Users will need to compile SYCL source code with C++ compilers which support the following C++ features:

- All C++03 features, apart from Run Time Type Information
- Exception handling
- C++11 lambda functions
- C++11 variadic templates
- C++11 template aliases

- C++11 rvalue references
- C++11 `std::function`, `std::string` and `std::vector`, although users can optionally define and use their own versions of these classes, which SYCL can use via template aliases.

SYCL programs are explicitly parallel and expose the full heterogeneous parallelism of the underlying machine model of OpenCL. This includes exposing the data-parallelism, multiple execution devices and multiple memory storage spaces of OpenCL. However, SYCL adds on top of OpenCL a higher level of abstraction allowing developers to hide much of the complexity from the source code, when a developer so chooses.

A SYCL program is logically split into host code and kernels. Host code is standard C++ code, as provided by whatever C++ compiler the developer chooses to use for the host code. The kernels are C++ *functors* (see C++ documentation for an explanation of functors) or C++11 lambda functions which have been designated to be compiled as SYCL kernels. SYCL will also accept OpenCL `cl_kernel` objects.

SYCL programs target heterogeneous systems. The kernels may be compiled and optimized for multiple different processor architectures with very different binary representations. In SYCL, kernels are contained within command group functors, which include all of the data movement/mapping/copying required to correctly execute the kernel. A command group functor takes as a parameter a command group `handler` class object which provides the interface for kernel invocations and then the functor itself can be submitted to a queue. The command group functor is executed on the host in order to add all the specified commands to the specified queue. Kernels can only access shared data via `accessor` objects constructed within the command group functor by passing the command group `handler` as a parameter.

2.5.1 Basic data parallel kernels

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

2.5.2 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 `cl::sycl::parallel_for` function with a `cl::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 `barrier` function on an `nd_item` object. 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.

2.5.3 Hierarchical data parallel kernels

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 `cl::sycl::parallel_for` the user calls `cl::sycl::parallel_for_work_group` with a `cl::sycl::range` value representing the number of work-groups to launch and optionally a second `cl::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 workgroup 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.

2.5.4 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 `cl::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, like any other OpenCL entity, be executed in task-parallel fashion.

2.5.5 Synchronization

In SYCL, synchronization can be either global or local within a work-group. The SYCL implementation may need to provide extra synchronization commands and host-side synchronization in order to enable synchronization across OpenCL contexts, but this is handled internally within the SYCL host runtime.

Synchronization between work-items in a single work-group is achieved using a work-group barrier. This matches the OpenCL C behaviour. All the work-items of a work-group must execute the barrier before any are allowed to continue execution beyond the barrier. Note that the work-group barrier must be encountered by all work-items of a work-group executing the kernel or by none at all. There is no mechanism for synchronization between work-groups. In SYCL, workgroup barriers are exposed through a method on the `cl::sycl::nd_item` class, which is only available inside kernels that are executed over workgroups. This ensures that developers can only use workgroup barriers inside workgroups.

Synchronization points in SYCL are exposed through the following operations:

- *Buffer destruction:* The destructors for `cl::sycl::buffer` and `cl::sycl::image` objects wait for all enqueued work on those objects to complete. If the objects were constructed with attached host memory, then the destructor copies the data back to host memory before returning. 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`.
- *Accessor construction:* 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 submitted to any queue will wait for the accessor to be destroyed.
- *Command group enqueue:* The SYCL scheduler 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 `handler_event` are returned by the queue's submit function that contain event information related to the specific command group.

- *Interaction with OpenCL synchronization operations:* The user can obtain OpenCL events from command groups, images and buffers which will enable the user to add barrier packets to their own queues to correctly synchronize for buffer or image data dependencies.
- *Queue destruction:* The destructor for `cl::sycl::queue` objects waits for all commands executing on the queue to complete before the destructor returns.
- *Context destruction:* The destructor for `cl::sycl::context` objects waits for all commands executing on any queues in the context to complete before the destructor returns.
- *SYCL event objects:* SYCL provides `cl::sycl::event` objects which can be used for user synchronization. If synchronization is required between two different OpenCL contexts, then the SYCL runtime ensures that any extra host-based synchronization is added to enable the SYCL event objects to operate between contexts correctly.

2.5.6 Error handling

In SYCL, there are two types of error: synchronous errors that can be detected immediately, and asynchronous errors that can only be detected later. 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 user-supplied asynchronous error-handlers.

A `cl::sycl::context` can be constructed with a user-supplied asynchronous error handler. If a `cl::sycl::queue` is constructed without a user-supplied context, then the user can supply an asynchronous error handler for the queue, otherwise errors on that queue will be reported to its context error handler.

Asynchronous errors are not reported immediately as they occur. The asynchronous error handler for a context or queue is called with a `cl::sycl::exception_list` object, which contains a list of asynchronously-generated exception objects, either on destruction of the context or queue that the error handler is associated with, or via an explicit `wait_and_throw` method call on an associated queue. This style of asynchronous error handling is similar to that proposed for an upcoming revision of the C++ standard.

2.5.7 Scheduling of kernels and data movement

Within the command group scope, accessor objects specify what data the command group will read and write. When enqueueing a command group functor, the runtime ensures that synchronization operations are also enqueued. By these means it ensures that the reads and writes are semantically equivalent to an in-order execution. Different command groups may execute out-of-order relative to each other, as long as read and write dependencies are enforced.

A command group functor can be submitted either to a single queue to be executed on, or a secondary queue can be provided as well. If a command group functor 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 functor 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 functor 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.

A command group functor at construction takes a command group [handler](#) as a parameter and anything within that scope is immediately executed and has to get the handler object as a parameter. The intention is that a user will perform calls to SYCL functions, methods, destructors and constructors inside that scope. These calls will be non-blocking on the host, but enqueue operations to the queue the command group is submitted at. All user functions within the command group scope will be called on the host as the command group functor is executed, but any runtime SYCL operations will be queued.

The scheduler must treat command groups atomically. So if two threads simultaneously enqueue two command groups onto the same queue, then each command group must be added to the queue as an atomic operation. The order of two simultaneously enqueued command groups relative to each other is undefined but the constituent commands must not interleave.

Command group functors are scheduled to enforce the ordering semantics of operations on memory objects (both buffers and images). These ordering rules apply regardless of whether the command groups are enqueued in the same context, queue, device or platform. Therefore, a SYCL implementation may need to produce extra synchronization operations between contexts, platforms, devices and queues using OpenCL constructs such as user events. How this is achieved is implementation defined. An implementation is free to re-order or parallelize command groups in queues as long as the ordering semantics on memory objects are not violated.

The ordering semantics on memory objects are:

1. The ordering rules apply based on the totality of accessors constructed in the command group. The order in which accessors are constructed within the command group is not relevant. If multiple accessors in the same command group operate on the same memory object, then the command group's access to that memory object is the union of the access permissions of the accessors.
2. Accessors can be created to operate on *sub-buffers*. A buffer may be overlaid with any number of sub-buffers. 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.
3. If a command group has any accessor with *discard* access to a memory object, then the scheduler does not need to preserve the previous contents of the memory object when scheduling the command group.
4. All other accessors must preserve normal read-write ordering and data access. This means the scheduler must ensure that a command group that reads a memory object must first copy or map onto the device the data that might be read. Reads must follow writes to memory objects or overlapping sub-buffers.
5. 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 scheduler to maintain multiple read-only copies of the data on multiple devices.

In OpenCL, there are in-order queues and out-of-order queues. In SYCL, the SYCL queue is always out-of-order, irrespective of the underlying OpenCL queues, maintaining the in-order execution semantics. There are no guarantees, however, that all the applications can be executed in an out-of-order scheduling scheme so the execution of command group functors cannot be guaranteed to be out-of-order or executed independently in parallel by the final system. The latter is implementation defined.

It is worth noting that a SYCL queue does not necessarily map to only one OpenCL queue, however, the OpenCL queue that is given when interacting with the SYCL queue will retain any synchronization information is needed for synchronization with any other OpenCL queues spawned by the system.

An OpenCL implementation can require different queues for different devices and contexts. The synchronization required to ensure order between commands in different queues varies according to whether the queues have shared contexts. A SYCL implementation must determine the required synchronization to ensure the above ordering rules above are enforced.

SYCL provides *host accessors*. These accessors give temporary access to data in buffers on the host, outside the command group scope. Host accessors are the only kinds of accessors that can be created outside command groups. Creation of a host accessor is a blocking operation: all command groups that read or write data in the buffer or image that the host accessor targets must have completed before the host thread will continue. All data being written in an enqueued command group to the buffer or image must be completed and written to the associated host memory before the host accessor constructor returns. Any subsequently enqueued command group that accesses overlapping data in the buffer or image of the host accessor will block and not start execution until the host accessor (and any copies) has been destroyed. This approach guarantees that there is no concurrent access to a memory object between the host thread and any SYCL device.

If a user creates a SYCL buffer, image or accessor from an OpenCL object, then the SYCL runtime will correctly manage synchronization and copying of data between the OpenCL memory object for the lifetime of the SYCL buffer, image or accessor constructed from it. If a user makes use of the underlying OpenCL memory object at the same time as a SYCL buffer, image or accessor is live, then the behaviour is undefined.

2.5.8 Managing object lifetimes

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

All OpenCL 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 OpenCL context) for the lifetime of their application to initialize and release the OpenCL context safely.

When an OpenCL object that is encapsulated in a SYCL object is copied in C++, then the underlying OpenCL object is not duplicated, but its OpenCL reference count is incremented. When the original or copied SYCL object is destroyed, then the OpenCL reference count is decremented.

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 will be non-blocking on destruction. In the case

where host memory is shared between the user application and the SYCL runtime, then the reference counter of the `shared_ptr` is determining 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.

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 on host memory before it starts processing. The host accessor does not necessarily copy back to the same host memory as the one initially given by the user.
- Memory object destruction, in the case where copies back to host memory have to be done
- Queue destruction, as all enqueued kernels need to finish executing first.
- Context destruction, as all enqueued kernels need to finish executing first.
- Device destruction, as all enqueued kernels need to finish executing first.
- Platform destruction, as all enqueued kernels need to finish executing first.

2.5.9 Device discovery and selection

A user specifies which queue to submit a command group functor on 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 selector causes the SYCL runtime to perform device discovery. No device discovery is performed until a SYCL 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 both OpenCL devices and platforms as well as a SYCL host platform and SYCL host device. The host device allows queue creation and running of kernels, but does not support OpenCL-specific features. It is an error for a user to request an underlying OpenCL device for the SYCL host device.

2.5.10 Interfacing with OpenCL

All SYCL objects which encapsulate an OpenCL object (such as contexts or queues) can be constructed from the OpenCL object. The constructor takes one argument, the OpenCL object, and performs an OpenCL `retain` operation on the OpenCL object to increase its reference count. The destructor for the SYCL object performs an OpenCL `release` operation on the OpenCL object. The copy construction semantics of the SYCL object ensure that each new SYCL copy of the object also does an OpenCL `retain` on the underlying object.

To obtain the underlying OpenCL object from a SYCL object, there is a `get` method on all relevant SYCL objects. The `get` method returns the underlying OpenCL object and also performs a `retain` operation on the object. It is the user's responsibility to `release` the OpenCL object when the user has finished with it.

SYCL images and buffers are treated differently in that SYCL image and buffer objects do not refer to an OpenCL context and may reference multiple underlying OpenCL image or buffer objects as well as host allocations. It is the accessors to the image and buffer objects that refer to an actual OpenCL context. Accessors provide synchronization in place of the events that the OpenCL runtime would use directly. Therefore, obtaining OpenCL `cl_mem`

objects from SYCL is achieved via special accessor classes which can return OpenCL `cl_mem` and `cl_event` objects. SYCL memory objects can be constructed from `cl_mem` objects, but the SYCL system is free to copy from the OpenCL memory object into another memory object or host memory, to achieve normal SYCL semantics, for as long as the SYCL memory object is live. No SYCL object is guaranteed to have only one underlying OpenCL object created, however, every SYCL object is required to have an OpenCL object which an OpenCL program can interface with and have all synchronization points refer to it.

2.6 Anatomy of a SYCL application

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

```
1 #include <CL/sycl.hpp>
2 #include <iostream>
3
4 int main() {
5     using namespace cl::sycl;
6
7     int data[1024]; // initialize data to be worked on
8
9     // By including all the SYCL work in a {} block, we ensure
10    // all SYCL tasks must complete before exiting the block
11    {
12        // create a queue to enqueue work to
13        queue myQueue;
14
15        // wrap our data variable in a buffer
16        buffer<int, 1> resultBuf(data, range<1>(1024));
17
18        // create a command_group to issue commands to the queue
19        myQueue.submit([&(handler& cgh) {
20            // request access to the buffer
21            auto writeResult = resultBuf.get_access<access::write>(cgh);
22
23            // enqueue a prallel_for task
24            cgh.parallel_for<class simple_test>(range<1>(1024), [=](id<1> idx) {
25                writeResult[idx] = static_cast<int>(idx[0]);
26            }); // end of the kernel function
27        }); // end of our commands for this queue
28    } // end of scope, so we wait for the queued work to complete
29
30    // print result
31    for (int i = 0; i < 1024; i++)
32        std::cout<<"data["<<i<<" = "<<data[i]<<std::endl;
33
34    return 0;
35 }
```

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

A SYCL application has 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 23 to 25. The *command group scope* specifies a unit of work which will comprise of a *kernel function* and *accessors*. In this example *command group scope* is defined by lines 18 to 26. 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.

A *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, by the `operator()` function of a function object or by the binary `cl_kernel` entity generated from an OpenCL C string. Each instance of the *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 is templated with a class, in this case called `class simple_test`. This class is used only as a name to enable the kernel (compiled with a device compiler) and the host code (possibly compiled with a different host compiler) to be linked. This is required because C++ lambda functions have no name that a linker could use to link the kernel to the host code.

The `parallel_for` method 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 *kernel function* will be executed over the given range from 0 to 1023. A *kernel function* can only be defined within a *command group scope*. Command group scope is the syntactic scope wrapped by the construction of a command group functor object as seen on line 18. The command group functor takes as a parameter a command group `handler` which is a runtime constructed object. The command group functor contains all the the operations to be enqueued on the queue this functor will be submitted to. In this case the constructor used for `myQueue` on line 13 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 *kernel function* must be contained within a *buffer* or *image*. We construct a buffer on line 16. Access to the buffer is controlled via an *accessor* which is constructed on line 21 through the `get_access` method 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 *kernel functions*. In this example the *accessor* is used to write to the data buffer on line 25. All *buffers* must be constructed in the *application scope*, whereas all *accessors* must be constructed in the *command group scope*.

2.7 Memory objects

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 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, float), vector data type, or a user-defined structure. In SYCL, a buffer object is a templated type (`cl::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 OpenCL implementation. Images are encapsulated in the `cl::sycl::image` type, which is 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 can be converted within a kernel into C++ pointer types, or the `cl::sycl::global_ptr`, `cl::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 stored in the same format as it is accessed by the kernel, but in the case of an image object the data format used to store the image elements may not be the same as the data format used inside the kernel. 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 methods to read from an image convert an image element from the format it is stored into a 4-component vector. Similarly, the SYCL accessor methods provided to write to an image convert the image element from a 4-component vector to the appropriate image format specified such as 4 8-bit elements, for example.

Memory objects, both buffers and images, may have one or more underlying OpenCL `cl_mem` objects. 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.

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, C++ allocator classes are used that can be the pre-defined ones or user-defined. For better control of synchronization between a SYCL and a non SYCL application that share data `shared_ptr` and `mutex` classes are used. In the case where the user would not like the host side to block on destruction of buffers or images, as the data given to the buffers are for initialization only, the `unique_ptr` class can be used instead of a raw pointer to data.

2.8 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 layer provides a set of C++ templates and classes which provide the programming model to the user. It enables the creation of queues, buffers and images, as well as access to some underlying OpenCL features 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 a SYCL-specific host device to run kernels on.
- *SYCL Device Compiler(s)*: 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.

2.9 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.

2.9.1 Building a SYCL program

A SYCL program runs on a *host* and one or more OpenCL 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 host.

The design of SYCL enables a single SYCL source file to be passed to multiple, different compilers. 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 OpenCL 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:

1. *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 compiler and the 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 (potentially a macro) 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.

2.9.2 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 runtime to load the kernel and pass it to an OpenCL runtime. This requires the kernel to have a globally-visible name to enable an association between the kernel invocation and the kernel itself. The association is achieved using a *kernel name*, which is a C++ typename.

For a functor, the kernel name can be the same type as the functor itself, as long as the functor type is globally accessible. For a lambda function, there is no globally-visible name, so the user must provide one. In SYCL, the name is provided as a template parameter to the kernel invocation, e.g. `parallel_for<kernelName>`.

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. For details please refer to 5.2. The user can also extract OpenCL `cl_kernel` and `cl_program` objects for kernels by providing the typename of the kernel.

2.10 Language restrictions in kernels

The SYCL kernels are executed on SYCL devices and all of the functions called from a SYCL kernel is going to be compiled for the device by a SYCL device compiler. Due to restrictions of the OpenCL 1.2 runtime and OpenCL 1.2 capable devices, there are certain restrictions for SYCL kernels. Those restrictions can be summarized as the kernels cannot include RTTI information, exception classes, recursive code, virtual functions or make use of C++ libraries that are not compiled for the device. For more details on language restrictions please refer to 5.3.

SYCL kernels use parameters that are captured by value in the command group scope described in 3.5.2 or are passed from the host to the device using the data management runtime classes of `cl::sycl::accessors`. Sharing data structures between host and device code imposes certain restrictions, such as you can only use user defined classes that are *C++11 standard layout* classes for the data structures and in general, no pointers initialized for host can be used on the device, the only way of passing pointers to a kernel is through the usage of `cl::sycl::accessor` class, which supports the `cl::sycl::buffer` and `cl::sycl::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 3.5.7.

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`. It is the responsibility of the SYCL device compiler to ensure that the sizes of these types match the sizes on the host, to enable data of these types to be shared between host and device.

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 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.

2.10.1 SYCL Linker

In SYCL only offline linking is supported for SYCL and OpenCL programs and libraries. In the case of linking C++ functions and methods to a SYCL application, where the definitions and declarations are not available in the same translation unit of the compiler, then the macro `SYCL_EXTERNAL` has to be provided. Any OpenCL C function included in a pre-built OpenCL library can be defined as `extern "C"` function and the 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 data types defined in 3.73.

2.10.2 Functions and datatypes available in kernels

Inside kernels, the functions and datatypes available are restricted by the underlying capabilities of OpenCL devices. All OpenCL C features are provided by C++ classes and functions, which are available on host and device.

2.11 Execution of kernels on the SYCL host device

SYCL enables kernels to run on either the host device or on OpenCL devices. When kernels run on an OpenCL device, then the features and behaviour of that execution follows the OpenCL specification, otherwise they follow the behaviour specified for the SYCL host device.

Any kernel enqueued to a host queue executes on the host device according to the same rules as the OpenCL devices.

Kernel math library functions on the host must conform to OpenCL math precision requirements.

The range of image formats supported by the host device is implementation- defined, but must match the minimum requirements of the OpenCL specification.

Some of the OpenCL extensions and optional features may be available on a SYCL host device, but since these are optional features and vendor specific extensions, the user must query the host device to determine availability. A SYCL implementer must state what OpenCL device features are available on their host device implementation.

The synchronization and data movement that occurs when a kernel is executed on the host may be implemented in a variety of ways on top of OpenCL. The actual mechanism is implementation-defined.

2.12 Example SYCL application

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

```
1  #include <CL/sycl.hpp>
2  #include <iostream>
3
4  using namespace cl::sycl;
5
6  // Size of the matrices
7  const size_t N = 2000;
8  const size_t M = 3000;
9
10 int main() {
11 {
12 // By including all the SYCL work in a {} block, we ensure
13 // all SYCL tasks must complete before exiting the block
14
15
16 // Open scope to make lifetime of queue clear
```



```

17 {
18     // Create a queue to work on
19     queue myQueue;
20
21     // Create some 2D buffers of float for our matrices
22     buffer<float, 2> a({N, M});
23     buffer<float, 2> b({N, M});
24     buffer<float, 2> c({N, M});
25
26     // Launch a first asynchronous kernel to initialize a
27     myQueue.submit( [&](handler& cgh) {
28         // The kernel write a, so get a write accessor on it
29         auto A = a.get_access<access::write>(cgh);
30
31         // Enqueue a parallel kernel iterating on a N*M 2D iteration space
32         cgh.parallel_for<class init_a>(
33             {N, M}, [=](id<2> index) { A[index] = index[0] * 2 + index[1]; });
34     });
35
36     // Launch an asynchronous kernel to initialize b
37     myQueue.submit( [&](handler& cgh) {
38         // The kernel write b, so get a write accessor on it
39         auto B = b.get_access<access::write>(cgh);
40         /* From the access pattern above, the SYCL runtime detect this
41            command_group is independant from the first one and can be
42            scheduled independently */
43
44         // Enqueue a parallel kernel iterating on a N*M 2D iteration space
45         cgh.parallel_for<class init_b>({N, M}, [=](id<2> index) {
46             B[index] = index[0] * 2014 + index[1] * 42;
47         });
48     });
49
50     // Launch an asynchronous kernel to compute matrix addition c = a + b
51     myQueue.submit( [&](handler& cgh) {
52         // In the kernel a and b are read, but c is written
53         auto A = a.get_access<access::read>(cgh);
54         auto B = b.get_access<access::read>(cgh);
55         auto C = c.get_access<access::write>(cgh);
56         // From these accessors, the SYCL runtime will ensure that when
57         // this kernel is run, the kernels computing a and b completed
58
59         // Enqueue a parallel kernel iterating on a N*M 2D iteration space
60         cgh.parallel_for<class matrix_add>(
61             {N, M}, [=](id<2> index) { C[index] = A[index] + B[index]; });
62     });
63
64     /* Ask an access to read c from the host-side. The SYCL runtime
65        ensures that c is ready when the accessor is returned */
66     auto C = c.get_access<access::read, access::host_buffer>();
67     std::cout << std::endl << "Result:" << std::endl;
68     for (size_t i = 0; i < N; i++) {
69         for (size_t j = 0; j < M; j++) {
70             // Compare the result to the analytic value
71             if (C[i][j] != i * (2 + 2014) + j * (1 + 42)) {

```

```
72         std::cout << "Wrong value " << C[i][j] << " on element " << i << " "
73         << j << std::endl;
74         exit(-1);
75     }
76 }
77 }
78 } /* End scope of myQueue, this wait for any remaining operations on the
79 queue to complete */
80
81 std::cout << "Good computation!" << std::endl;
82 return 0;
83 }
```

3. SYCL Programming Interface

The SYCL programming interface provides a C++ abstraction to OpenCL 1.2 functionality and feature set. This section describes all the available classes and interfaces of SYCL, focusing on the C++ interface of the underlying runtime. In this section, we are defining all the classes and methods for the SYCL API, which are available for SYCL host and OpenCL devices. This section also describes the synchronization rules and OpenCL API interoperability rules which guarantee that all the methods, including constructors, of the SYCL classes are thread safe.

It is assumed that the OpenCL API is also available to the developer at the same time as SYCL.

3.1 Header files and namespaces

SYCL provides one standard header file: "`CL/sycl.hpp`", which needs to be included in every SYCL program.

All SYCL classes, constants, types and functions are defined within the `cl::sycl` namespace.

3.2 C++ Standard library classes required for the interface

The SYCL programming interfaces make extensive use of vectors, strings and function objects to carry information. Moreover, smart pointer and mutex classes, allow extending the SYCL programming interface in terms of host data management a SYCL will default to using the STL string, vector, function, mutex and smart pointer classes, unless defined otherwise. These types are exposed internally as `cl::sycl::vector_class`, `cl::sycl::string_class`, `cl::sycl::function_class`, `cl::sycl::mutex_class`, `cl::sycl::unique_ptr_class`, `cl::sycl::shared_ptr_class` and `cl::sycl::weak_ptr_class`.

It is possible to disable the STL versions of these classes when required. A common reason for doing this is to specify a custom allocator to move memory management under the control of the SYCL user. This is achieved by defining `CL_SYCL_NO_STD_VECTOR`, `CL_SYCL_NO_STD_STRING`, `CL_SYCL_NO_STD_FUNCTION`, `CL_SYCL_NO_STD_MUTEX`, `CL_SYCL_NO_UNIQUE_PTR`, `CL_SYCL_NO_SHARED_PTR`, `CL_SYCL_NO_WEAK_PTR`, respectively, before including "`CL/sycl.hpp`", and by replacing the template aliases in the `cl::sycl` namespace as necessary.

```
1 namespace cl {
2 namespace sycl {
3     #define CL_SYCL_NO_STD_VECTOR
4     #include <vector>
5     template < class T, class Alloc = std::allocator<T> >
6     using vector_class = std::vector<T, Alloc>;
7
8     #define CL_SYCL_NO_STD_STRING
9     #include <string>
```

```

10     using string_class = std::string;
11
12     #define CL_SYCL_NO_STD_FUNCTION
13     #include <functional>
14     using function_class = std::function
15
16     #define CL_SYCL_NO_STD_MUTEX
17     #include <functional>
18     using mutex_class = std::mutex
19
20     #define CL_SYCL_NO_STD_UNIQUE_PTR
21     #include <functional>
22     using unique_ptr_class = std::unique_ptr
23
24     #define CL_SYCL_NO_STD_SHARED_PTR
25     #include <functional>
26     using shared_ptr_class = std::shared_ptr
27
28     #define CL_SYCL_NO_STD_WEAK_PTR
29     #include <functional>
30     using weak_ptr_class = std::weak_ptr
31
32     #include <cl/sycl.hpp>
33 } // sycl
34 } // cl

```

3.3 SYCL runtime classes

3.3.1 Device selection class

The class `device_selector` is a functor which enables the SYCL runtime to choose the best device based on heuristics specified by the user, or by one of the built-in device selectors. The built-in device selectors are listed in Table 3.3. The `device_selector` constructors and methods are described in tables 3.1 and 3.2.

```

1 namespace cl {
2 namespace sycl {
3 class device_selector {
4 public:
5     device_selector();
6
7     device_selector(const device_selector &selector);
8
9     virtual ~device_selector();
10
11     device select_device() const;
12
13     virtual int operator()(const device &device) const = 0;
14 };
15 } // namespace sycl
16 } // namespace cl

```

Constructors	Description
<code>device_selector()</code>	Default device selector constructor for the abstract class.
<code>device_selector(const device_selector &selector)</code>	Copy constructor.
End of table	

Table 3.1: Constructors of the `device_selector` class

Methods	Description
<code>device select_device()const</code>	Returns a selected device using the functor operator defined in sub-classes <code>operator()(const device &device)</code> .
<code>virtual int operator()(const device &device)const</code>	This pure virtual operator allows the customization of device selection. It defines the behavior of the <code>device_selector</code> functor called by the SYCL runtime on device selection. It returns a “score” for each device in the system and the highest rated device will be used by the SYCL runtime.
End of table	

Table 3.2: Methods for the `device_selector` class

`operator()` is an abstract method which returns a “score” per-device. At the stage where the SYCL runtime selects a device, the system will go through all the available devices in the system and choose the one with the highest score as computed by the current device selection class. If a device has a negative score it will never be chosen. While OpenCL devices may or may not be available, the SYCL host device is always available, so the developer is able to choose the SYCL host device as a fall-back device. Selection of the SYCL host device will allow execution of CPU-compiled versions of kernels scheduled on queues created against that device.

The system also provides built-in device selectors, including selectors which choose a device based on the default behavior of the system. An important note is that the system is not required to have global state and its behavior is defined by the platforms the developer chooses to target.

The *default_selector* is the selector that incorporates the *default* behavior of the system, and it is implicitly used by the system for the creation of the queue when no other *device_selector* or underlying OpenCL identifier is provided. The method the default selector uses to rank and select devices is implementation-defined. The *default_selector* will choose the SYCL host device if there are no OpenCL devices available.

SYCL device selectors	Description
<code>default_selector</code>	Devices selected by heuristics of the system. If no OpenCL device is found then it defaults to the SYCL host device.
<code>gpu_selector</code>	Select devices according to device type <code>info::device::device_type::gpu</code> from all the available OpenCL devices. If no OpenCL GPU device is found the selector fails.
Continued on next page	

Table 3.3: Standard device selectors included with all SYCL implementations.

SYCL device selectors	Description
cpu_selector	Select devices according to device type <code>info::device::device_type::cpu</code> from all the available devices and heuristics. If no OpenCL CPU device is found the selector fails.
host_selector	Selects the SYCL host CPU device that does not require an OpenCL runtime.
End of table	

Table 3.3: Standard device selectors included with all SYCL implementations.

3.3.2 Platform class

The `platform` class represents a SYCL platform: a collection of related SYCL supported devices. Each platform may be either an OpenCL platform, or it may be the SYCL host platform, containing only the SYCL host device. The host platform may be useful when no OpenCL platform is available or during the development process, especially for debugging. The platform class offers a selection of static methods to obtain information about the platforms available at runtime. The SYCL host platform reports itself as a valid SYCL platform. The constructors and methods of the `platform` class are listed in Tables 3.4 and 3.5.

```

1 namespace cl {
2 namespace sycl {
3 namespace info {
4 enum class device_type : unsigned int {
5     cpu,
6     gpu,
7     accelerator,
8     custom,
9     defaults,
10    host,
11    all
12 }
13 } // info
14
15 class platform {
16 public:
17     platform();
18
19     explicit platform(cl_platform_id platformID);
20
21     explicit platform(device_selector &devSelector);
22
23     platform(const platform &rhs);
24
25     platform &operator=(const platform &rhs);
26
27     ~platform();
28
29     //The OpenCL cl_platform_id or nullptr for SYCL host.

```

```

30  cl_platform_id get() const;
31
32  // Returns all the available OpenCL platforms and the SYCL host platform
33  static vector_class<platform> get_platforms() const;
34
35  //Returns the devices available in this platform
36  vector_class<device> get_devices(
37      info::device_type = info::device_type::all) const;
38
39  //Returns the corresponding descriptor information for all SYCL platforms
40  //(OpenCL and host)
41  template <info::platform param>
42  typename info::param_traits<info::platform, param>::type get_info()
43      const;
44
45  //Returns the available extensions for all SYCL platforms( OpenCL and host)
46  bool has_extension(string_class extension) const;
47
48  //True if the platform is host
49  bool is_host() const;
50 };
51 } // namespace sycl
52 } // namespace cl

```

The SYCL host platform is not an OpenCL platform. The *get()* method will return nullptr *cl_platform_id*. The SYCL host platform will be included in the output of the static function *get_platforms*.

Constructors	Description
<code>platform()</code>	Default constructor for platform. It constructs a platform object to encapsulate the device returned by the default <i>device_selector</i> . Returns errors via the SYCL exception class.
<code>explicit platform(cl_platform_id platformId)</code>	Construct a platform object from an OpenCL platform id. Returns errors via the SYCL exception class.
<code>explicit platform(device_selector &devSelector)</code>	Construct a platform object from the device returned by a device selector of the user's choice. Returns errors via the SYCL exception class.
<code>platform(const platform &rhs)</code>	Copy constructor.
<code>platform &operator=(const platform &rhs)</code>	Assignment operator.
End of table	

Table 3.4: Constructors of `platform` class

The default constructor will create an instance of the platform class where the underlying platform will be the SYCL host platform by default.

Methods	Description
<code>cl_platform_id get ()const</code>	Returns the <code>cl_platform_id</code> of the underlying OpenCL platform. If the platform is not a valid OpenCL platform, for example if it is the SYCL host, a <code>nullptr</code> will be returned.
<code>static vector_class<platform> get_platforms ()const</code>	Returns all available platforms in the system.
<code>vector_class<device> get_devices(info::device_type = info::device_type::all)const</code>	Returns all the available devices for this platform, of type <code>device_type</code> , which is defaulted to <code>info::device_type::all</code>
<code>template <info::platform param> typename info::param_traits<info::platform, param>::type get_info()const</code>	Queries OpenCL information for the underlying <code>cl_platform</code> .
<code>bool has_extension(string_class extension)const</code>	Specifies whether a specific extension is supported on the platform.
<code>bool is_host()const</code>	Returns true if this is a SYCL host platform.
End of table	

Table 3.5: Methods of `platform` class

3.3.2.1 Platform information descriptors

A SYCL platform can be queried for all of the following information using the `get_info` function. All SYCL contexts have valid devices for them, including the SYCL host device. The available information is in table 3.6. The interface of all available `nplatform` descriptors in the appendix C.1.

Platform Descriptors	Return type	Description
<code>info::platform::profile</code>	<code>string_class</code>	Returns the profile name supported by the implementation. Can be either <code>FULL_PROFILE</code> or <code>EMBEDDED_PROFILE</code> .
<code>info::platform::version</code>	<code>string_class</code>	OpenCL software driver version string in the form <code>major_number.minor_number</code>
<code>info::platform::name</code>	<code>string_class</code>	Name of the platform.
<code>info::platform::vendor</code>	<code>string_class</code>	String provided by the platform vendor.
<code>info::platform::extensions</code>	<code>string_class</code>	A space-separated list of extension names supported by the platform.
End of table		

Table 3.6: Platform information descriptors.

3.3.3 Context class

The `context` class encapsulates an OpenCL context, which is implicitly created and the lifetime of the context instance defines the lifetime of the underlying OpenCL context instance. On destruction `clReleaseContext` is called. The default context is the SYCL host context containing only the SYCL host device.

The constructors and methods of the `context` class are listed in Tables 3.7 and 3.8.

3.3.3.1 Context interface

```
1 namespace cl {
2 namespace sycl {
3 class context {
4 public:
5     context();
6
7     explicit context(async_handler asyncHandler = nullptr);
8
9     context(cl_context clContext, async_handler asyncHandler = nullptr);
10
11     context(const device_selector &deviceSelector,
12             info::gl_context_interop interopFlag,
13             async_handler asyncHandler = nullptr);
14
15     context(const device &dev, info::gl_context_interop interopFlag,
16             async_handler asyncHandler = nullptr);
17
18     context(const platform &plt, info::gl_context_interop interopFlag,
19             async_handler asyncHandler = nullptr);
20
21     context(vector_class<device> deviceList, info::gl_context_interop interopFlag,
22             async_handler asyncHandler = nullptr);
23
24     context(const context &rhs);
25
26     context &operator=(const context &rhs);
27
28     ~context();
29
30     cl_context get() const;
31
32     bool is_host() const;
33
34     platform get_platform();
35
36     vector_class<device> get_devices() const;
37
38     template <info::context param>
39     typename param_traits<info::context, param>::type get_info() const;
40 };
41 } // namespace sycl
42 } // namespace cl
```

Constructors	Description
<code>context ()</code>	Default constructor that chooses the context according the heuristics of the <i>default-selector</i> . Returns synchronous errors via the SYCL exception class.
<code>explicit context(async_handler asyncHandler = nullptr)</code>	Constructs a context object for SYCL host using an <code>async_handler</code> for handling asynchronous errors.
<code>context(const device_selector &deviceSelector, info::gl_context_interop interopFlag, async_handler asyncHandler = nullptr)</code>	Constructs a context object using a <code>device_selector</code> object. The context is constructed with a single device retrieved from the <code>device_selector</code> object provided. Returns synchronous errors via the SYCL exception class and asynchronous errors are handled via the <code>async_handler</code> , if provided.
<code>context(const device &dev, info::gl_context_interop interopFlag, async_handler asyncHandler = nullptr)</code>	Constructs a context object using a device object. Returns synchronous errors via the SYCL exception class and asynchronous errors are handled via the <code>async_handler</code> , if provided.
<code>context(const platform &plt, info::gl_context_interop interopFlag, async_handler asyncHandler = nullptr)</code>	Constructs a context object using a platform object. Returns synchronous errors via the SYCL exception class and asynchronous errors are handled via the <code>async_handler</code> , if provided.
<code>context(vector_class<device> deviceList, info::gl_context_interop interopFlag, async_handler asyncHandler = nullptr)</code>	Constructs a context object using a <code>vector_class</code> of device objects. Returns synchronous errors via the SYCL exception class and asynchronous errors are handled via the <code>async_handler</code> , if provided.
<code>context (cl_context clContext, async_handler asyncHandler = nullptr)</code>	Context constructor, where the underlying OpenCL context is given as a parameter. The constructor executes a <code>retain</code> on the <i>cl_context</i> . Returns synchronous errors via the SYCL exception class and asynchronous errors are handled via the <code>async_handler</code> , if provided.
<code>context(const context &rhs)</code>	Constructs a context object from another context object and retains the <code>cl_context</code> object if the context is not SYCL host.
End of table	

Table 3.7: Constructors of the `context` class.

Methods	Description
<code>cl_context get ()const</code>	Returns the underlying <i>cl_context</i> object, after retaining the <code>cl_context</code> . Retains a reference to the returned <code>cl_context</code> object. Caller should release it when finished.

Methods	Description
<code>bool is_host ()const</code>	Specifies whether the context is in SYCL Host Execution Mode.
<code>template <info::context param> typename param_traits <info::context, param>::type get_info ()const</code>	Queries OpenCL information for the underlying <code>cl_context</code> .
<code>platform get_platform()</code>	Returns the SYCL platform that the context is initialized for.
<code>vector_class<device> get_devices()const</code>	Returns the set of devices that are part of this context.

Table 3.8: Methods of context class

3.3.3.2 Context information descriptors

A SYCL context can be queried for all of the following information using the `get_info` function. All SYCL contexts have valid devices for them, including the SYCL host device. The available information is in table 3.9. The interface of all available context descriptors in the appendix C.2. UPDATED: changed `info::context::context_-reference_count` to `info::context::reference_count` for consistency)

Context Descriptors	Return type	Description
<code>info::context::reference_count</code>	<code>cl_uint</code>	Return the context reference count.
<code>info::context::num_devices</code>	<code>cl_uint</code>	Return the number of devices in context.
<code>info::context::devices</code>	<code>vector_class< cl_device_id></code>	Return the list of devices in context.
<code>info::context::gl_interop</code>	<code>info:: gl_context_interop</code>	Boolean value which specifies whether the context is used for OpenCL/OpenGL interoperability according to the OpenCL 1.2 extensions specification document [2].
		End of table

Table 3.9: Context information descriptors

On construction of a context, it is possible to supply an asynchronous error handler function object. If supplied, then asynchronous errors can be reported to the error handler. Asynchronous errors are only reported to the user when a queue attached to the context is destroyed or as its `wait_and_throw()` method called.

3.3.4 Device class

The SYCL `device` class encapsulates a particular SYCL device against on which kernels may be executed. The SYCL device may be an OpenCL device or it may be a SYCL host device representing the host CPU. In the OpenCL device case it should have valid `cl_device_id` and `cl_platform_id` available. The `cl_device_id` for the SYCL host device is not going to be available through the OpenCL interface, as it is not an OpenCL device. In the case where the SYCL device is constructed from an existing `cl_device_id` the system will call `clRetainDevice`. On destruction the runtime will call `clReleaseDevice`. It is the user's responsibility to make sure that the device

object with `cl_device_id` is a valid object during the lifetime of the `device` class.

3.3.4.1 Device interface

The constructors and methods of the `device` class are listed in Tables 3.10 and 3.11.

```
1 namespace cl {
2 namespace sycl {
3 class device {
4 public:
5     device();
6
7     explicit device(cl_device_id deviceId);
8
9     explicit device(device_selector &deviceSelector);
10
11     device(const device &rhs);
12
13     device &operator=(const device &rhs);
14
15     ~device();
16
17     // The OpenCL cl_platform_id or nullptr for SYCL host.
18     cl_device_id get() const;
19
20     bool is_host() const;
21
22     bool is_cpu() const;
23
24     bool is_gpu() const;
25
26     bool is_accelerator() const;
27
28     platform get_platform() const;
29
30     // Returns all the available OpenCL devices and the SYCL host device
31     static vector_class<device> get_devices(
32         info::device_type deviceType = info::device_type::all);
33
34     template <info::device param>
35     typename info::param_traits<info::device, param>::type
36     get_info<info::device>() const;
37
38     bool has_extension(string_class extension) const;
39
40     vector_class<device> create_sub_devices(
41         info::device_partition_type partitionType,
42         info::device_partition_property partitionProperty,
43         info::device_affinity_domain affinityDomain) const;
44
45
46 };
47 } // namespace sycl
```

```
48 } // namespace cl
```

The default constructor will create an instance of the SYCL host device.

The developer can partition existing devices through the `create_sub_devices` API. More documentation on this is in the OpenCL 1.2 specification [1, sec. 4.3]. It is valid to construct a SYCL device directly from an OpenCL sub-device.

Information about the SYCL device may be queried through the `get_info` method. The developer can also query the device instance for the `cl_device_id` which will be `nullptr` if the SYCL device is the host device. The `get_platform` method will return the corresponding `platform` object.

To facilitate the different options for SYCL devices, there are methods that check the type of device. The method `is_host()` returns true if the device is actually the host. In the case where an OpenCL device has been initialized through this API, the methods `is_cpu()`, `is_gpu()` and `is_accelerator()` return true if the OpenCL device is CPU, GPU or an accelerator.

Constructors	Description
<code>device ()</code>	Default constructor for the device. It chooses a device using <code>host_selector</code> . Returns errors via C++ exception class.
<code>explicit device (device_selector &deviceSelector)</code>	Constructs a device class instance using the <code>device_selector</code> provided. Returns errors via C++ exception class.
<code>explicit device (cl_device_id deviceId)</code>	Constructs a device class instance using <code>cl_device_id</code> of the OpenCL device. Returns synchronous errors via the SYCL exception class. Retains a reference to the OpenCL device and if this device was an OpenCL sub-device the device should be released by the caller when it is no longer needed.
<code>device (const device &rhs)</code>	Copy constructor. Returns synchronous errors via the SYCL exception class.
<code>device &operator=(const device &rhs)</code>	Assignment constructor. Returns synchronous errors via the SYCL exception class.
End of table	

Table 3.10: Constructors of the `device` class

Methods	Description
<code>cl_device_id get ()const</code>	Returns the <code>cl_device_id</code> of the underlying OpenCL platform. Returns synchronous errors via the SYCL exception class. Retains a reference to the returned <code>cl_device_id</code> object. Caller should release it when finished. In the case where this is the SYCL host device it will return a <code>nullptr</code> .
Continued on next page	

Table 3.11: Methods of the `device` class

Methods	Description
<code>platform get_platform ()const</code>	Returns the platform of device. Returns synchronous errors via the SYCL exception class.
<code>bool is_host ()const</code>	Returns true if the device is a SYCL host device.
<code>bool is_cpu ()const</code>	Returns true if the device is an OpenCL CPU device.
<code>bool is_gpu ()const</code>	Returns true if the device is an OpenCL GPU device.
<code>bool is_accelerator ()const</code>	Returns true if the device is an OpenCL accelerator device.
<code>template <info::device param> typename info::param_traits <info::device, param>::type get_info ()const</code>	Queries the device for OpenCL <i>info::device_info</i> . Returns synchronous errors via the SYCL exception class.
<code>bool has_extension (string_class extension)const</code>	Specifies whether a specific extension is supported on the device.
<code>vector_class<device> create_sub_devices (info::device_partition_type partitionType, info::device_partition_property partitionProperty, info::device_affinity_domain affinityDomain) const</code>	Partitions the device into sub devices based upon the properties provided. Returns synchronous errors via SYCL exception classes.
<code>static vector_class<device> get_devices (info::device_type deviceType = info::device_type::all)</code>	Returns a list of all available devices. Returns synchronous errors via SYCL exception classes.
End of table	

Table 3.11: Methods of the `device` class

3.3.4.2 Device information descriptors

A SYCL device can be queried for all of the following information using the `get_info` function. All SYCL devices have valid queries for them, including the SYCL host device. The available information is in table 3.12. The interface of all available device descriptors in the appendix C.3.

Device Descriptors	Return type	Description
<code>info::device::device_type</code>	<code>info::device_type</code>	The SYCL device type. Currently supported values are: <code>cpu</code> , <code>gpu</code> , <code>accelerator</code> , <code>defaults</code> , <code>custom</code> , <code>host</code> , <code>all</code> .
<code>info::device::vendor_id</code>	<code>cl_uint</code>	A unique SYCL device vendor identifier. An example of a unique device identifier could be the PCIe ID. The SYCL host device has to report a valid vendor id.
Continued on next page		

Table 3.12: Device information descriptors.

Device Descriptors	Return type	Description
<code>info::device::max_compute_units</code>	<code>cl_uint</code>	The number of parallel compute units on the SYCL device. A work-group executes on a single compute unit. The minimum value is 1.
<code>info::device::max_work_item_dimensions</code>	<code>cl_uint</code>	Maximum dimensions that specify the global and local work-item IDs used by the data parallel execution model. The minimum value is 3 for devices that are not of type <code>info::device_type::custom</code> .
<code>info::device::max_work_item_sizes</code>	<code>id<3></code>	Maximum number of work-items that can be specified in each dimension of the work-group to the <code>nd_range</code> . The minimum value is (1, 1, 1) for devices that are not of type <code>info::device_type::custom</code> .
<code>info::device::max_work_group_size</code>	<code>size_t</code>	Maximum number of work-items in a work-group executing a kernel on a single compute unit, using the data parallel execution model. The minimum value is 1.
<code>info::device::preferred_vector_width_char</code> <code>info::device::preferred_vector_width_short</code> <code>info::device::preferred_vector_width_int</code> <code>info::device::preferred_vector_width_long</code> <code>info::device::preferred_vector_width_float</code> <code>info::device::preferred_vector_width_double</code> <code>info::device::preferred_vector_width_half</code>	<code>cl_uint</code>	Preferred native vector width size for builtin scalar types that can be put into vectors. The vector width is defined as the number of scalar elements that can be stored in the vector. If double precision is not supported, <code>info::device_preferred_width_double</code> must return 0. If the <code>cl_khr_fp16</code> extension is not supported, <code>info::device::preferred_vector_width_half</code> must return 0.
<code>info::device::native_vector_width_char</code> <code>info::device::native_vector_width_short</code> <code>info::device::native_vector_width_int</code> <code>info::device::native_vector_width_long</code> <code>info::device::native_vector_width_float</code> <code>info::device::native_vector_width_double</code> <code>info::device::native_vector_width_half</code>	<code>cl_uint</code>	Returns the native ISA vector width. The vector width is defined as the number of scalar elements that can be stored in the vector. If double precision is not supported, <code>info::device::native_vector_width_double</code> must return 0. If the <code>cl_khr_fp16</code> extension is not supported, <code>info::device::native_vector_width_half</code> must return 0.
<code>info::device::max_clock_frequency</code>	<code>cl_uint</code>	Maximum configured clock frequency of the device in MHz.
<code>info::device::address_bits</code>	<code>cl_uint</code>	The default compute device address space size specified as an unsigned integer value in bits. Currently supported values are 32 or 64 bits.

Continued on next page

Table 3.12: Device information descriptors.

Device Descriptors	Return type	Description
<code>info::device::max_mem_alloc_size</code>	<code>cl_long</code>	Max size of memory object allocation in bytes. The minimum value is $\max(1/4\text{th of } \text{info::device::global_mem_size}, 128*1024*1024)$ for devices that are not of type <code>info::device::type::custom</code> .
<code>info::device::image_support</code>	<code>cl_bool</code>	Is 1 if images are supported by the SYCL device and 0 otherwise.
<code>info::device::max_read_image_args</code>	<code>cl_uint</code>	Max number of simultaneous image objects that can be read by a kernel. The minimum value is 128 if <code>info::device::image_support</code> is true.
<code>info::device::max_write_image_args</code>	<code>cl_uint</code>	Max number of simultaneous image objects that can be written to by a kernel. The minimum value is 8 if <code>info::device::image_support</code> is true.
<code>info::device::image2d_max_width</code>	<code>size_t</code>	Max width of 2D image or 1D image not created from a buffer object in pixels. The minimum value is 8192 if <code>info::device::image_support</code> is true.
<code>info::device::image_2d_max_height</code>	<code>size_t</code>	Max height of 2D image in pixels. The minimum value is 8192 if <code>info::device::image_support</code> is true.
<code>info::device::image3d_max_width</code>	<code>size_t</code>	Max width of 3D image in pixels. The minimum value is 2048 if <code>info::device::image_support</code> is true.
<code>info::device::image3d_max_height</code>	<code>size_t</code>	Max height of 3D image in pixels. The minimum value is 2048 if <code>info::device::image_support</code> is true.
<code>info::device::image3d_max_depth</code>	<code>size_t</code>	Max depth of 3D image in pixels. The minimum value is 2048 if <code>info::device::image_support</code> is true.
<code>info::device::image_max_buffer_size</code>	<code>size_t</code>	Max number of pixels for a 1D image created from a buffer object. The minimum value is 65536 if <code>info::device::image_support</code> is true.
<code>info::device::image_max_array_size</code>	<code>size_t</code>	Max number of images in a 1D or 2D image array. The minimum value is 2048 if <code>info::device::image_support</code> is true.
<code>info::device::max_samplers</code>	<code>cl_uint</code>	Maximum number of samplers that can be used in a kernel. The minimum value is 16 if <code>info::device::image_support</code> is true.
<code>info::device::max_parameter_size</code>	<code>size_t</code>	Max size in bytes of the arguments that can be passed to a kernel. The minimum value is 1024 for devices that are not of type <code>info::device::type::custom</code> . For this minimum value, only a maximum of 128 arguments can be passed to a kernel.

Continued on next page

Table 3.12: Device information descriptors.

Device Descriptors	Return type	Description
<code>info::device::mem_base_addr_align</code>	<code>cl_uint</code>	The minimum value is the size (in bits) of the largest SYCL built-in data type supported by the device is <code>longlong16</code> for devices that are not of type <code>info::device_type::custom</code> .
<code>info::device::single_fp_config</code>	<code>info::device_fp_config</code>	Describes single precision floating-point capability of the device. This is a bit-field that describes one or more of the following values: <ul style="list-style-type: none"> • <code>info::fp_config::denorm</code> : denorms are supported • <code>info::fp_config::inf_nan</code> : INF and quiet NaNs are supported. • <code>info::fp_config::round_to_nearest</code>: round to nearest even rounding mode supported • <code>info::fp_config::round_to_zero</code> : round to zero rounding mode supported • <code>info::fp_config::round_to_inf</code> : round to positive and negative infinity rounding modes supported • <code>info::fp_config::fma</code> : IEEE754-2008 fused multiply add is supported. • <code>info::fp_config::correctly_rounded_divide_sqrt</code> : divide and sqrt are correctly rounded as defined by the IEEE754 specification. • <code>info::fp_config::soft_float</code> : Basic floating-point operations (such as addition, subtraction, multiplication) are implemented in software. The mandated minimum floating-point capability for devices that are not of type <code>info::device_type::custom</code> is: <code>info::fp_config::round_to_nearest</code> <code>info::fp_config::inf_nan</code>

Continued on next page

Table 3.12: Device information descriptors.

Device Descriptors	Return type	Description
<code>info::device::double_fp_config</code>	<code>info::device_fp_config</code>	<p>Describes double precision floating-point capability of the SYCL device. This is a bit-field that describes one or more of the following values:</p> <ul style="list-style-type: none"> • <code>info::fp_config::denorm</code>: denorms are supported • <code>info::fp_config::inf_nan</code>: INF and NaNs are supported. • <code>info::fp_config::round_to_nearest</code>: round to nearest even rounding mode supported. • <code>info::fp_config::round_to_zero</code>: round to zero rounding mode supported. • <code>info::fp_config::round_to_inf</code>: round to positive and negative infinity rounding modes supported. • <code>info::fp_config::fma</code>: IEEE754-2008 fused multiply-add is supported. • <code>info::fp_config::soft_float</code>: Basic floating-point operations (such as addition, subtraction, multiplication) are implemented in software. <p>Double precision is an optional feature so the mandated minimum double precision floating-point capability is 0. If double precision is supported by the device, then the minimum double precision floating-point capability must be: <code>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 info::fp_config::denorm</code>.</p>
<code>info::device::global_mem_cache_type</code>	<code>info::device_mem_cache_type</code>	Type of global memory cache supported. Valid values are: none, read_only_cache, write_only_cache.
<code>info::device::global_mem_cacheline_size</code>	<code>cl_uint</code>	Size of global memory cache line in bytes.
<code>info::device::global_mem_cache_size</code>	<code>cl_ulong</code>	Size of global memory cache in bytes.
<code>info::device::global_mem_size</code>	<code>cl_ulong</code>	Size of global device memory in bytes.
<code>info::device::max_constant_buffer_size</code>	<code>cl_ulong</code>	Max size in bytes of a constant buffer allocation. The minimum value is 64 KB for devices that are not of type <code>info::device_type::custom</code> .
<code>info::device::max_constant_args</code>	<code>cl_uint</code>	Max number of constant arguments declared in a kernel. The minimum value is 8 for devices that are not of type <code>info::device_type::custom</code> .

Continued on next page

Table 3.12: Device information descriptors.

Device Descriptors	Return type	Description
<code>info::device::local_mem_type</code>	<code>info::local_mem_type</code>	Type of local memory supported. This can be set to <code>info::local_mem_type::local</code> implying dedicated local memory storage such as SRAM, or <code>info::local_mem_type::global</code> . For custom devices, <code>info::local_mem_type::none</code> can also be returned indicating no local memory support.
<code>info::device::local_mem_size</code>	<code>cl_ulong</code>	Size of local memory arena in bytes. The minimum value is 32 KB for devices that are not of type <code>info::device_type::custom</code> .
<code>info::device::error_correction_support</code>	<code>cl_bool</code>	Is true if the device implements error correction for all accesses to compute device memory (global and constant). Is false if the device does not implement such error correction.
<code>info::device::host_unified_memory</code>	<code>cl_bool</code>	Is true if the device and the host have a unified memory subsystem and is false otherwise.
<code>info::device::profiling_timer_resolution</code>	<code>size_t</code>	Describes the resolution of device timer. This is measured in nanoseconds.
<code>info::device::is_endian_little</code>	<code>cl_bool</code>	Is true if the SYCL device is a little endian device and false otherwise.
<code>info::device::is_available</code>	<code>cl_bool</code>	Is true if the device is available and false if the device is not available.
<code>info::device::is_compiler_available</code>	<code>cl_bool</code>	Is false if the implementation does not have a compiler available to compile the program source. An OpenCL device that conforms to the OpenCL Embedded Profile may not have an online compiler available.
<code>info::device::is_linker_available</code>	<code>cl_bool</code>	Is false if the implementation does not have a linker available. An OpenCL device that conforms to the OpenCL Embedded Profile may not have a linker available. However, it needs to be true if <code>info::device::is_compiler_available</code> is true.
<code>info::device::execution_capabilities</code>	<code>info::device_execution_capabilities</code>	Describes the execution capabilities of the device. This is a bit-field that describes one or more of the following values: <ul style="list-style-type: none"> <code>info::device_execution_capabilities::exec_kernel</code> : The OpenCL device can execute OpenCL kernels. <code>info::device_execution_capabilities::exec_native_kernel</code> : The OpenCL device can execute native kernels. The mandated minimum capability is: <code>info::device_execution_capabilities::exec_kernel</code>.

Continued on next page

Table 3.12: Device information descriptors.

Device Descriptors	Return type	Description
<code>info::device::queue_properties</code>	<code>vector_class<info::device_queue_properties></code>	Describes the command-queue properties supported by the device. This is a bit-field that describes one or more of the following values: <ul style="list-style-type: none"> <code>info::queue_properties::out_of_order</code> <code>info::queue_properties::profiling</code> The mandated capability is to enable profiling.
<code>info::device::built_in_kernels</code>	<code>string_class</code>	A semi-colon separated list of built-in kernels supported by the device. An empty string is returned if no built-in kernels are supported by the device.
<code>info::device::platform</code>	<code>cl_platform_id</code>	The platform associated with this device.
<code>info::device::name</code>	<code>string_class</code>	Device name string
<code>info::device::vendor</code>	<code>string_class</code>	Vendor name string.
<code>info::device::version</code>	<code>string_class</code>	OpenCL software driver version string in the form <code>major_number.minor_number</code>
<code>info::device::profile</code>	<code>string_class</code>	OpenCL profile string. Returns the profile name supported by the device. The profile name returned can be one of the following strings: <ul style="list-style-type: none"> <code>FULL_PROFILE</code> : if the device supports the OpenCL specification (functionality defined as part of the core specification and does not require any extensions to be supported). <code>EMBEDDED_PROFILE</code> - if the device supports the OpenCL embedded profile.
<code>info::device::version</code>	<code>string_class</code>	OpenCL version string. Returns the OpenCL version supported by the device. This version string has the following format: <code>OpenCL<space><major_version.minor_version><space><vendor-specific-information></code> The <code>major_version.minor_version</code> value returned will be 1.2.

Continued on next page

Table 3.12: Device information descriptors.

Device Descriptors	Return type	Description
info::device::opencl_c_version	string_class	<p>OpenCL C version string. Returns the highest OpenCL C version supported by the compiler for this device that is not of type info::device_type::custom. This version string has the following format: OpenCL<space>C<space><major_version.minor_version><space><vendor-specific-information></p> <p>The major_version.minor_version value returned must be 1.2 if info::device::version is OpenCL 1.2. The major_version.minor_version value returned must be 1.1 if info::device::version is OpenCL 1.1. The major_version.minor_version value returned can be 1.0 or 1.1 if info::device::version is OpenCL 1.0.</p>

Continued on next page

Table 3.12: Device information descriptors.

Device Descriptors	Return type	Description
<code>info::device::extensions</code>	<code>string_class</code>	<p>Returns a space separated list of extension names (the extension names themselves do not contain any spaces) supported by the device. The list of extension names returned can be vendor supported extension names and one or more of the following Khronos approved extension names:</p> <ul style="list-style-type: none"> • <code>cl_khr_int64_base_atomics</code> • <code>cl_khr_int64_extended_atomics</code> • <code>cl_khr_3d_image_writes</code> • <code>cl_khr_fp16</code> • <code>cl_khr_gl_sharing</code> • <code>cl_khr_gl_event</code> • <code>cl_khr_d3d10_sharing</code> • <code>cl_khr_dx9_media_sharing</code> • <code>cl_khr_d3d11_sharing</code> • <code>cl_khr_depth_images</code> • <code>cl_khr_gl_depth_images</code> • <code>cl_khr_gl_msaa_sharing</code> • <code>cl_khr_image2d_from_buffer</code> • <code>cl_khr_initialize_memory</code> • <code>cl_khr_context_abort</code> • <code>cl_khr_spir</code> <p>The following approved Khronos extension names must be returned by all device that support OpenCL C 1.2:</p> <ul style="list-style-type: none"> • <code>cl_khr_global_int32_base_atomics</code> • <code>cl_khr_global_int32_extended_atomics</code> • <code>cl_khr_local_int32_base_atomics</code> • <code>cl_khr_local_int32_extended_atomics</code> • <code>cl_khr_byte_addressable_store</code> • <code>cl_khr_fp64</code> (for backward compatibility if double precision is supported) <p>Please refer to the OpenCL 1.2 Extension Specification for a detailed description of these extensions.</p>
<code>info::device::printf_buffer_size</code>	<code>size_t</code>	<p>Maximum size of the internal buffer that holds the output of <code>printf</code> calls from a kernel. The minimum value for the full profile is 1 MB.</p>

Continued on next page

Table 3.12: Device information descriptors.

Device Descriptors	Return type	Description
<code>info::device::preferred_interop_user_sync</code>	<code>cl_bool</code>	Is true if the devices preference is for the user to be responsible for synchronization, when sharing memory objects between OpenCL and other APIs such as DirectX, false if the device/implementation has a performant path for performing synchronization of memory object shared between OpenCL and other APIs such as DirectX.
<code>info::device::parent_device</code>	<code>cl_device_id</code>	Returns the <code>cl_device_id</code> of the parent device to which this sub-device belongs. If device is a root-level device, or the host device a NULL value is returned.
<code>info::device::partition_max_sub_devices</code>	<code>cl_uint</code>	Returns the maximum number of subdevices that can be created when a device is partitioned. The value returned cannot exceed <code>info::device::device_max_compute_units</code>
<code>info::device::partition_properties</code>	<code>vector_class</code> < <code>info::device_partition_property</code> >	Returns the list of partition types supported by device. The is an array of <code>info::device_partition_property</code> values drawn from the following list: <ul style="list-style-type: none"> <code>info::device_partition_property::partition_equally</code> <code>info::device_partition_property::partition_by_counts</code> <code>info::device_partition_by_affinity_domain</code> If the device cannot be partitioned (i.e. there is no partitioning scheme supported by the device that will return at least two subdevices), a value of 0 will be returned.
<code>info::device::partition_affinity_domain</code>	<code>info::device_affinity_domain</code>	Returns the list of supported affinity domains for partitioning the device using <code>info::device_affinity_domain</code> . This is a bit-field that describes one or more of the following values: <ul style="list-style-type: none"> <code>info::device_affinity_domain::unsupported</code> <code>info::device_affinity_domain::numa</code> <code>info::device_affinity_domain::L4_cache</code> <code>info::device_affinity_domain::L3_cache</code> <code>info::device_affinity_domain::L2_cache</code> <code>info::device_affinity_domain::L1_cache</code> <code>info::device_affinity_domain::next_partitionable</code>

Continued on next page

Table 3.12: Device information descriptors.

Device Descriptors	Return type	Description
<code>info::device::partition_type</code>	<code>vector< info::device_partition_property ></code>	<p>Returns the properties argument specified when creating sub devices, if device is a subdevice. In the case where the properties argument to creating sub devices is <code>info::device_partition_property::partition_by_affinity_domain</code> <code>info::device_partition_property::partition_affinity_domain_next_partitionable</code>, the affinity domain used to perform the partition will be returned. This can be one of the following values:</p> <ul style="list-style-type: none"> • <code>info::device_affinity_domain::unsupported</code> • <code>info::device_affinity_domain::numa</code> • <code>info::device_affinity_domain::L4_cache</code> • <code>info::device_affinity_domain::L3_cache</code> • <code>info::device_affinity_domain::L2_cache</code> • <code>info::device_affinity_domain::L1_cache</code> <p>Otherwise the implementation may either return a <code>param_value_size_ret</code> of 0 i.e. there is no partition type associated with device or can return a property value of 0 (where 0 is used to terminate the partition property list) in the memory that <code>param_value</code> points to.</p>
<code>info::device::reference_count</code>	<code>cl_uint</code>	Returns the device reference count. If the device is a root-level device, a reference count of one is returned.
End of table		

Table 3.12: Device information descriptors.

3.3.5 Queue class

The class `queue` is a SYCL encapsulation of an OpenCL `cl_command_queue`. Can be constructed from a `cl_command_queue`. The destructor waits for all execution on the queue to end and then passes any exceptions that occurred asynchronously on the queue to the asynchronous error handler, if provided to the constructor, before calling `clReleaseCommandQueue`.

3.3.5.1 Queue interface

The constructors and methods of the `Queue` class are listed in Tables 3.13 and 3.14.

```

1 namespace cl {
2 namespace sycl {
3 class queue {
4 public:
5     explicit queue(async_handler asyncHandler = nullptr);

```



```

6
7   queue(const device_selector &deviceSelector,
8         async_handler asyncHandler = nullptr);
9
10  queue(const context &syclContext, const device_selector &deviceSelector,
11        async_handler asyncHandler = nullptr);
12
13  queue(const context &syclContext, const device &syclDevice,
14        async_handler asyncHandler = nullptr);
15
16  queue(const context &syclContext, const device &syclDevice, info
17        : queue_profiling profilingFlag, async_handler asyncHandler = nullptr);
18
19  queue(const device &syclDevice, async_handler asyncHandler = nullptr);
20
21  queue(cl_command_queue clQueue, async_handler asyncHandler = nullptr);
22
23  queue(const queue &syclQueue);
24
25  ~queue();
26
27  bool is_host();
28
29  context get_context();
30
31  cl_command_queue get();
32
33  device get_device();
34
35  context get_context();
36
37  template <info::queue param>
38  typename param_traits<info::queue, param>::type get_info() const;
39
40  template <typename T>
41  handler_event submit(T cgf);
42
43  template <typename T>
44  handler_event submit(T cgf, queue &secondaryQueue);
45
46  void wait();
47
48  void wait_and_throw();
49
50  void throw_asynchronous();
51 };
52 } // namespace sycl
53 } // namespace cl

```

Constructors	Description
<code>explicit queue (async_handler asyncHandler = nullptr)</code>	This constructor creates a SYCL queue from an OpenCL queue. At construction it does a retain on the queue memory object. Retains a reference to the <code>cl_command_queue</code> object. Caller should release the passed <code>cl_command_queue</code> object when it is no longer needed. Returns synchronous errors regarding the creation of the queue and reports asynchronous errors via the <code>async_handler</code> callback function in conjunction with the synchronization and throw methods.
<code>queue (const device_selector &selector, async_handler asyncHandler = nullptr)</code>	Creates a queue for the device provided by the <code>device_selector</code> . If no device is selected, an error is reported. Returns synchronous errors regarding the creation of the queue and reports asynchronous errors via the <code>async_handler</code> callback function if and only if there is an <code>async_handler</code> provided.
<code>queue (const device &syclDevice, async_handler asyncHandler = nullptr)</code>	A queue is created for <code>syclDevice</code> . Returns asynchronous errors via the <code>async_handler</code> callback function.
<code>queue (context &syclContext, device_selector &deviceSelector async_handler asyncHandler = nullptr)</code>	This constructor chooses a device based on the provided <code>device_selector</code> , which needs to be in the given context. If no device is selected, an error is reported. Returns synchronous errors regarding the creation of the queue. If and only if there is an <code>asyncHandler</code> provided, it reports asynchronous errors via the <code>async_handler</code> callback function in conjunction with the synchronization and throw methods.
<code>queue (const context &syclContext, const device &syclDevice async_handler asyncHandler = nullptr)</code>	Creates a command queue using <code>clCreateCommandQueue</code> from a context and a device. Returns synchronous errors regarding the creation of the queue. If and only if there is an <code>asyncHandler</code> provided, it reports asynchronous errors via the <code>async_handler</code> callback function in conjunction with the synchronization and throw methods.

Continued on next page

Table 3.13: Constructors of the `queue` class.

Constructors	Description
<code>queue (const context &syclContext, const device &syclDevice, info::queue_profiling profilingFlag, async_handler asyncHandler = nullptr)</code>	Creates a command queue using <code>clCreateCommandQueue</code> from a context and a device. It enables profiling on the queue if the <code>profilingFlag</code> is set to true. Returns synchronous errors regarding the creation of the queue. If and only if there is an <code>asyncHandler</code> provided, it reports asynchronous errors via the <code>async_handler</code> callback function in conjunction with the <code>synchronization</code> and <code>throw</code> methods.
<code>queue (cl_command_queue clQueue, async_handler asyncHandler = nullptr)</code>	This constructor creates a SYCL queue from an OpenCL queue. At construction it does a retain on the queue memory object. Returns synchronous errors regarding the creation of the queue. If and only if there is an <code>asyncHandler</code> provided, it reports asynchronous errors via the <code>async_handler</code> callback function in conjunction with the <code>synchronization</code> and <code>throw</code> methods.
<code>queue (queue &syclQueue)</code>	Copy constructor
End of table	

Table 3.13: Constructors of the `queue` class.

Methods	Description
<code>cl_command_queue get()</code>	Returns the underlying OpenCL command queue after doing a retain. This memory object is expected to be released by the developer. Retains a reference to the returned <code>cl_command_queue</code> object. Caller should release it when finished. If the queue is a SYCL host queue then a <code>nullptr</code> will be returned.
<code>context get_context ()</code>	Returns the SYCL queue's context. Reports errors using SYCL exception classes.
<code>device get_device ()</code>	Returns the SYCL device the queue is associated with. Reports errors using SYCL exception classes.
<code>bool is_host()</code>	Returns whether the queue is executing on a SYCL host device.
<code>void wait ()</code>	Performs a blocking wait for the completion all enqueued tasks in the queue. Synchronous errors will be reported through SYCL exceptions.
Continued on next page	

Table 3.14: Methods for class `queue`

Methods	Description
<code>void wait_and_throw ()</code>	Performs a blocking wait for the completion all enqueued tasks in the queue. Synchronous errors will be reported via SYCL exceptions. Asynchronous errors will be passed to the <code>async_handler</code> passed to the queue on construction. If no <code>async_handler</code> was provided then asynchronous exceptions will be lost.
<code>void throw_asynchronous ()</code>	Checks to see if any asynchronous errors have been produced by the queue and if so reports them by passing them to the <code>async_handler</code> passed to the queue on construction. If no <code>async_handler</code> was provided then asynchronous exceptions will be lost.
<code>template<info::queue param> typename param_traits <info::queue, param>::type get_info ()const</code>	Queries the platform for <code>cl_command-queue.info</code>
<code>template <typename T> handler_event submit(T cgf)</code>	Submit a command group functor to the queue, in order to be scheduled for execution on the device.
<code>template <typename T> handler_event submit(T cgf, queue & secondaryQueue)</code>	Submit a command group functor to the queue, in order to be scheduled for execution on the device. On kernel error, this command group functor, then it is scheduled for execution on the secondary queue. Returns a command group functor event, which is corresponds to the queue the command group functor is being enqueued on.
End of table	

Table 3.14: Methods for class queue

3.3.5.2 Queue information descriptors

A SYCL queue can be queried for all of the following information using the `get_info` function. All SYCL queues have valid queries for them, including the SYCL host queue. The available information is in table 3.15. The interface of all available device descriptors in the appendix C.4.

Queue Descriptors	Return type	Description
<code>info::queue::context</code>	<code>cl_context</code>	Return the context specified when the command-queue is created.
<code>info::queue::device</code>	<code>cl_device_id</code>	Return the device specified when the command-queue is created.
<code>info::queue::reference_count</code>	<code>cl_uint</code>	Return the command-queue reference count.
Continued on next page		

Table 3.15: Queue information descriptors

Queue Descriptors	Return type	Description
<code>info::queue::properties</code>	<code>info::queue_profiling</code>	Return the currently specified properties for the command-queue. These properties are specified by the <code>properties</code> argument in the queue constructor.
		End of table

Table 3.15: Queue information descriptors

3.3.5.3 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 enqueueing work onto a queue. Such errors are returned through exceptions.
- **Asynchronous errors** are those that are produced through callback functions only. These will be stored within the queue's context until they are dispatched to the context's asynchronous error handler. If a queue is constructed with a user-supplied context, then it is this context's asynchronous error handler to which asynchronous errors are reported. If a queue is constructed without a user-supplied context, then the queue's constructor can be supplied with a queue-specific asynchronous error handler which will be used to construct the queue's context. To ensure that such errors are processed predictably in a known host thread these errors are only passed to the asynchronous error handler on request when either `wait_and_throw` is called or when `throw_asynchronous` is called. If no asynchronous error handler is passed to the queue or its context on construction, then such errors go unhandled, much as they would if no callback were passed to an OpenCL context.

3.3.6 Event class for OpenCL interoperability

An *event* in SYCL abstracts the `cl_event` objects in OpenCL. In OpenCL events' mechanism is comprised of low-level event objects that require from the developer to use them in order to synchronize memory transfers, enqueueing kernels and signaling barriers.

In SYCL, events are an abstraction of the OpenCL event objects, but they retain the features and functionality of the OpenCL event mechanism. They accommodate synchronization between different contexts, devices and platforms. It is the responsibility of the SYCL implementation to ensure that when SYCL events are used in OpenCL queues, the correct synchronization points are created to allow cross-platform or cross-device synchronization. Since data management and storage is handled by the SYCL runtime, the `event` class is used for providing the appropriate interface for OpenCL/SYCL interoperability. In the case where the SYCL objects contain OpenCL memory objects created outside the SYCL mechanism then events can be used in order to provide to the SYCL runtime the initial events it has to synchronize against. However, the events mechanism does not provide full interoperability during the SYCL code execution with OpenCL. Interoperability is achieved by using the synchronization rules with `buffer` and `image` class.

A SYCL event can be constructed from an OpenCL event or can return an OpenCL event. The constructors and methods of the Event class are listed in Tables 3.16 and 3.17.

```

1 namespace cl {
2 namespace sycl {
3 class event {
4 public:
5     event() = default;
6
7     explicit event(cl_event clEvent);
8
9     event(const event & rhs);
10
11     ~event();
12
13     cl_event get();
14
15     vector_class<event> get_wait_list();
16
17     void wait();
18
19     static void wait(const vector_class<event> &eventList);
20
21     void wait_and_throw();
22
23     static void wait_and_throw(const vector_class<event> &eventList);
24
25     template <info::event param>
26     typename param_traits<info::event, param>::type get_info() const;
27
28     template <info::event_profiling param>
29     typename param_traits<info::event_profiling, param>::type get_profiling_info() const;
30 };
31 } // namespace sycl
32 } // namespace cl

```

Constructors	Description
<code>event ()= default</code>	Default construct a null event object.
<code>explicit event (cl_event clEvent)</code>	Construct a SYCL event from a <code>cl_event</code> , only used with the SYCL/ OpenCL interoperability interface for buffers and images.
<code>event (const event & rhs)</code>	Construct a copy sharing the same underlying event. The underlying event will be reference counted.
End of table	

Table 3.16: Constructors for the `event` class

Methods	Description
<code>cl_event get()</code>	Return the underlying OpenCL event reference. Retains a reference to the returned <code>cl_event</code> object. Caller should release it when finished.
Continued on next page	

Table 3.17: Methods for the `event` class

Methods	Description
<code>vector_class<event> get_wait_list()</code>	Return the list of events that this event waits for in the dependence graph.
<code>void wait()</code>	Wait for the event and the command associated with it to complete.
<code>void wait_and_throw()</code>	Wait for the event and the command associated with it to complete. If any uncaught asynchronous errors occurred on the context (or contexts) that the event is waiting on executions from, then will also call that context's asynchronous error handler with those errors.
<code>static void wait(const vector_class<event> &eventList)</code>	Synchronously wait on a list of events.
<code>static void wait_and_throw(const vector_class<event> &eventList)</code>	Synchronously wait on a list of events. If any uncaught asynchronous errors occurred on the context (or contexts) that the events are waiting on executions from, then will also call those contexts' asynchronous error handlers with those errors.
<code>template <info::event param> typename param_traits <info::event, param>::type get_info ()const</code>	Queries OpenCL information for the underlying <i>cl.event</i> .
<code>template <info::event_profiling param> typename param_traits <info::event_profiling, param>::type get_profiling_info ()const</code>	Queries OpenCL profiling information for the underlying <i>cl.event</i> .
End of table	

Table 3.17: Methods for the `event` class

3.3.6.1 Event information and profiling descriptors

A SYCL event can be queried for all of the following information using the `get_info` function. The available information is in table 3.18. Profiling information available is in table 3.19. The interface of all available event and profiling descriptors in the appendix C.7.

Event Descriptors	Return type	Description
<code>info::event::command_type</code>	<code>cl_command_type</code>	Return the command associated with event. Can be one of the following values: <ul style="list-style-type: none"> • <code>CL_COMMAND_NDRANGE_KERNEL</code> • <code>CL_COMMAND_TASK</code> • <code>CL_COMMAND_NATIVE_KERNEL</code> • <code>CL_COMMAND_READ_BUFFER</code> • <code>CL_COMMAND_WRITE_BUFFER</code> • <code>CL_COMMAND_COPY_BUFFER</code> • <code>CL_COMMAND_READ_IMAGE</code> • <code>CL_COMMAND_WRITE_IMAGE</code> • <code>CL_COMMAND_COPY_IMAGE</code> • <code>CL_COMMAND_COPY_BUFFER_TO_IMAGE</code> • <code>CL_COMMAND_COPY_IMAGE_TO_BUFFER</code> • <code>CL_COMMAND_MAP_BUFFER</code> • <code>CL_COMMAND_MAP_IMAGE</code> • <code>CL_COMMAND_UNMAP_MEM_OBJECT</code> • <code>CL_COMMAND_MARKER</code> • <code>CL_COMMAND_ACQUIRE_GL_OBJECTS</code> • <code>CL_COMMAND_RELEASE_GL_OBJECTS</code> • <code>CL_COMMAND_READ_BUFFER_RECT</code> • <code>CL_COMMAND_WRITE_BUFFER_RECT</code> • <code>CL_COMMAND_COPY_BUFFER_RECT</code> • <code>CL_COMMAND_USER</code> • <code>CL_COMMAND_BARRIER</code> • <code>CL_COMMAND_MIGRATE_MEM_OBJECTS</code> • <code>CL_COMMAND_FILL_BUFFER</code> • <code>CL_COMMAND_FILL_IMAGE</code>
<code>info::event::command_execution_status</code>	<code>cl_int</code>	Execution status of the command underlying the event. One of: <code>CL_QUEUED</code> , <code>CL_SUBMITTED</code> , <code>CL_RUNNING</code> , <code>CL_COMPLETE</code> .
<code>info::event::reference_count</code>	<code>cl_uint</code>	Return the reference count of the event.
End of table		

Table 3.18: Event class information descriptors.

Event Profiling Descriptors	Return type	Description
<code>info::event_profiling::command_queued</code>	<code>cl_ulong</code>	64-bit value describing the device time in nanoseconds when the command was enqueued by the host.
<code>info::event_profiling::num_args</code>	<code>cl_ulong</code>	64-bit value describing the device time in nanoseconds when the command was submitted by the host to the device.
<code>info::event_profiling::reference_count</code>	<code>cl_ulong</code>	64-bit value describing the device time in nanoseconds when the command started executing on the device.
Continued on next page		

Table 3.19: Event class profiling information descriptors.

Event Profiling Descriptors	Return type	Description
<code>info::event_profiling::attributes</code>	<code>cl_ulong</code>	64-bit value describing the device time in nanoseconds when the command finished executing on the device.
		End of table

Table 3.19: Event class profiling information descriptors.

3.4 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 are different to OpenCL buffers and images in that they can be bound to more than one device or context and they get destroyed when they go out-of-scope. They also handle ownership of the data, while allowing exception handling for blocking and non-blocking data transfers. Accessors manage data transfers between host and all the devices in the system, as well as tracking data dependencies.

3.4.1 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 the other). 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 create data in 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 OpenCL contexts).

An allocator can always return the same host address if a user wants so.

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 mechanism (e.g. NUMA) and improves interoperability with STL-based libraries (e.g. Intel's TBB provides their own allocator).

3.4.1.1 Default Allocators

A default allocator is always defined by the implementation, and it is guaranteed to return non `nullptr` and new memory positions every time. 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
<code>buffer_allocator</code>	It is the default buffer allocator used by the runtime, when no allocator is defined by the user.
<code>image_allocator</code>	It is the default image allocator used by the runtime, when no allocator is defined by the user. The image allocator is required to be a byte-sized allocator, so the default type this allocator is typed to a type of size 1.
End of table	

Table 3.20: SYCL Default Allocators

See later 3.4.5 for detail on manual host-device synchronization.

3.4.1.2 Map Allocator

`map_allocator` is a class provided by the SYCL interface that can be matched with a specialized constructor of the buffer or image in order to provide the capability of 'mapping' the host data on any devices that the buffer uses.

This allocator always uses the same host address in order to create or map any device buffers, avoiding any copies from host data to host buffer and also uses the same host memory for all host accesses. A side effect of that, is that the host accessors will synchronize with the host memory and the runtime will handle that synchronization.

The host data address used for mapping the host data to the device data cannot be 'const', as the `map_allocator` will always use that address even for transferring data across devices of different contexts.

3.4.2 Buffers

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

A `cl::sycl::buffer` does not map to only one OpenCL buffer object and all OpenCL buffer memory objects are temporary for the use within a command group on a specific device. The only exception to this rule is when a buffer is constructed from a `cl_mem` object to interoperate with OpenCL. In the interop case the buffer will constitute a single `cl_mem` object and the ownership of the buffer `cl_mem` memory object remains at the OpenCL side. The SYCL buffer constructor and destructor use the existing retain and release mechanism available in OpenCL. Use of an interoperability buffer on a queue mapping to a context other than that in which the `cl_mem` was created is an error.

3.4.2.1 Buffer Interface

Buffer constructors are listed in Table 3.21 and methods in Table 3.22.

```
1 namespace cl{
2 namespace sycl {
3 template <typename T,
4         int dimensions,
5         typename AllocatorT = cl::sycl::buffer_allocator>
6 class buffer {
7 public:
8     using value_type = T;
9     using reference = value_type&;
10    using const_reference = const value_type&;
11
12    buffer(const range<dimensions> &bufferRange);
13
14    buffer(const T * hostData, const range<dimensions> & bufferRange);
15
16    buffer(T * hostData, const range<dimensions> & bufferRange);
17
18    buffer(shared_ptr_class<T> & hostData, const range<dimensions>
19           & bufferRange, cl::sycl::mutex_class * m = nullptr);
20
21    buffer(unique_ptr_class<void> && hostData,
22           const range<dimensions> & bufferRange);
23
24    buffer(buffer<T, dimensions, AllocatorT> b,
25           id<dimensions> & baseIndex,
26           range<dimensions> & subRange);
27
28    template <class InputIterator>
29    buffer<T, 1> (InputIterator first, InputIterator last);
30
31    buffer(cl_mem memObject,
32           queue & fromQueue, event availableEvent = {});
33
34    buffer(const buffer<T,dimensions,AllocatorT> &rhs);
35
36    buffer<T,dimensions,AllocatorT> &operator = (const
37           buffer<T,dimensions,AllocatorT> & rhs);
38
39    ~buffer();
40
41    size_t get_range();
42    size_t get_count();
43    size_t get_size();
44
45    template <access::mode mode,
46            access::target target=access::global_buffer>
47    accessor<T, dimensions, mode, target> get_access();
48
49    template<T>
50    set_final_data(weak_ptr_class<T> & finalData)
```

```

51 };
52 } // namespace sycl
53 } // namespace cl

```

Constructors	Description
<pre> template<typename T, int dimensions, typename AllocatorT = cl::sycl::buffer_allocator> > buffer (const range<dimensions> & bufferRange) </pre>	<p>Create a new buffer of the given size with storage managed by the SYCL runtime. The default behavior is to use the default host buffer allocator, in order to allow for host accesses. If the type of the buffer, has the const qualifier, then the default allocator will remove the qualifier to allow host access to the data.</p>
<pre> template<typename T, int dimensions, typename AllocatorT = cl::sycl::buffer_allocator> > buffer(const T* hostData, const range<dimensions> & bufferRange) </pre>	<p>Create a new buffer with associated host memory. <code>hostData</code> points to the storage and values used by the buffer and <code>range<dimensions></code> defines the size. The host address is const T, so the host accesses can be read-only. However, the <code>typename T</code> is not const so the device accesses can be both read and write accesses. Since, the <code>hostData</code> is const, this buffer is only initialized with this memory and there is no write after its destruction, unless there is another final data address given after construction of the buffer. The default value of the allocator is going to be the <code>cl::sycl::buffer_allocator</code> which will be of type T.</p>
<pre> template<typename T, int dimensions, typename AllocatorT= cl::sycl::buffer_allocator> buffer(T* hostData, const range<dimensions> & bufferRange) </pre>	<p>Create a new buffer with associated host memory. The memory is owned by the runtime during the lifetime of the object. Data is copied back to the host unless the user overrides the behavior using the <code>set_final_data</code> method. <code>hostData</code> points to the storage and values used by the buffer and <code>range<dimensions></code> defines the size.</p>
<pre> template<typename T, int dimensions, typename AllocatorT= cl::sycl::buffer_allocator> buffer(shared_ptr_class<T>& hostData, const range<dimensions> & bufferRange, cl::sycl ::mutex_class * m) </pre>	<p>Create a new buffer with associated memory, using the data in <code>codeinlinehostData</code>. The ownership of the <code>hostData</code> is shared between the runtime and the user. In order, to enable both the user application and the SYCL runtime to use the same pointer, an <code>cl::sycl::mutex_class</code> is used. The mutex <code>m</code> is locked by the runtime whenever the data is in use and unlocked otherwise. Data is synchronized with <code>hostData</code>, when the mutex is unlocked by the runtime.</p>

Continued on next page

Table 3.21: Constructors for the `buffer` class.

Constructors	Description
<pre>template<typename T, int dimensions, typename AllocatorT= cl::sycl::buffer_allocator> buffer(unique_ptr_class<void> && hostData, const range<dimensions> & bufferRange)</pre>	Create a new buffer which is initialized by hostData. The SYCL runtime receives full ownership of the hostData unique_ptr and there in effect there is no synchronization with the application code using hostData.
<pre>template<typename T, int dimensions =1, typename AllocatorT= cl::sycl::buffer_allocator> buffer(Iterator first, Iterator last)</pre>	Create a new allocated 1D buffer initialized from the given elements ranging from first up to one before last. The data is copied to an intermediate memory position by the runtime. Data is written back to the same iterator set if the iterator is not a const iterator.
<pre>template<typename T, int dimensions, typename AllocatorT= cl::sycl::buffer_allocator> buffer<T, dimensions, AllocatorT> (const buffer<T, dimensions, AllocatorT> & b)</pre>	Create a new buffer copy that shares the data than the original buffer. The system uses reference counting to deal with data lifetime. The destruction of a copy of a buffer does not trigger a copy back from the device.
<pre>template<typename T, int dimensions, typename AllocatorT= cl::sycl::buffer_allocator> buffer(buffer<T, dimensions, AllocatorT> &b, index<dimensions> & baseIndex, range<dimensions> & subRange)</pre>	Create a new sub-buffer without allocation to have separate accessors later. b is the buffer with the real data. baseIndex specifies the origin of the sub-buffer inside the buffer b. subRange specifies the size of the sub-buffer.
<pre>template<typename T, int dimensions, typename AllocatorT= cl::sycl::buffer_allocator> buffer (cl_mem memObject, queue & fromQueue, event availableEvent = {})</pre>	Create a buffer from an existing OpenCL memory object associated with a context after waiting for an event signaling the availability of the OpenCL data. memObject is the OpenCL memory object to use. fromQueue is the queue associated to the memory object. availableEvent specifies the event to wait for if non null. Note that a buffer created from a cl_mem object will only have one underlying cl_mem for the lifetime of the buffer and use on an incompatible queue constitutes an error.
<pre>buffer(const buffer<T, dimensions, AllocatorT> & rhs)</pre>	Copy constructor
End of table	

Table 3.21: Constructors for the `buffer` class.

Methods	Description
<pre>range<dimensions> get_range()</pre>	Return a range object representing the size of the buffer in terms of number of elements in each dimension as passed to the constructor.
<pre>size_t get_count()</pre>	Returns the total number of elements in the buffer. Equal to <code>get_range()[0] * ... * get_range()[dimensions-1]</code> .
Continued on next page	

Table 3.22: Methods for the `buffer` class.

Methods	Description
<code>size_t get_size()</code>	Returns the size of the buffer storage in bytes. Equal to <code>get_count()*sizeof(T)</code> .
<pre>template< access::mode mode, access::target target=access::global_buffer> accessor<T, dimensions, mode, target> get_access()</pre>	Returns a valid accessor to the buffer with the specified access mode and target. The value of target can be <code>access::global_buffer</code> , <code>access::constant_buffer</code> or <code>access::host_buffer</code> .
<pre>template<T> set_final_data(weak_ptr_class<T> & finalData)</pre>	<p>The <code>finalData</code> points to the host memory to which, the outcome of all the buffer processing is going to be copied to. This is the final pointer, which is going to be accessible after the destruction of the buffer and in the case where this is a valid pointer, the data are going to be copied to this host address.</p> <p><code>finalData</code> is different from the original host address, if the buffer was created associated with one. This is mainly to be used when a <code>shared_ptr</code> is given in the constructor and the output data will reside in a different location from the initialization data.</p> <p>It is defined as a <code>weak_ptr</code> referring to a <code>shared_ptr</code> that is not associated with the <code>cl::sycl::buffer</code>, and so the <code>cl::sycl::buffer</code> will have no ownership of <code>finalData</code>.</p>
End of table	

Table 3.22: Methods for the `buffer` class.

3.4.2.2 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

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 never blocks, 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 undefined.
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 undefined 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.

3. 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.

4. 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 accessed.

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.

However, if the `set_final_data()` function is used and host data pointer is valid for copying data back, then the buffer will block on destruction and copy data using the `weak_ptr` provided.

5. A buffer can be constructed using a `unique_ptr` to host data, which serve as an initialization point for the buffer. The ownership of the host data pointer is moved to the SYCL runtime, and those data will not be available after the destruction of the buffer.

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, since the original pointer is invalid.

However, if the `set_final_data()` function is used and host data pointer is valid for copying data back, then the buffer will block on destruction and copy data using the `weak_ptr` provided.

6. A buffer can be constructed using a `shared_ptr` to host data. This pointer is shared between the SYCL application and runtime. In order to allow synchronization between the the application and the runtime an `mutex` is used which will be locked by the runtime whenever the data are in use and unlock it when it no longer needs them.

The `shared_ptr` reference counting is used in order to prevent either from destroying the buffer host data prematurely. If the `shared_ptr` is deleted from the user application before the buffer destruction, the buffer can continue securely due to the fact that the pointer hasn't be destroyed yet, but will not copy data back to the host before destruction, as the application side has already deleted its copy.

There is no need to use the `set_final_data()` function in order to set the final data pointer if its is going to be the same as the `shared_ptr` used in construction of the buffer.

In the case where `set_final_data()` is used and a `weak_ptr` referencing to a valid `shared_ptr` on the SYCL application side, the buffer will block and copy data back to that host memory.

7. 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 will not block.
8. A buffer constructed from a `cl_mem` object creates a SYCL buffer that is initialized from a `cl_mem` object and may use the `cl_mem` object for the lifetime of the buffer. The destructor for this type of buffer will

block until all operations on the buffer have completed and then will (if necessary) copy all modified data back into the associated `cl_mem` object. The buffer will have a single `cl_mem` object and all operations will be performed on this underlying storage.

As a convenience for the user, any constructor that takes a range argument can instead be passed range values as 1, 2 or 3 arguments of type `size_t`.

A buffer object can also be copied, which just copies a reference to the buffer. The buffer objects use reference counting, so copying a buffer object increments a reference count on the underlying buffer. If after destruction, the reference count for the buffer is non-zero, then no further action is taken.

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, but which only have access to the range specified at time of construction of the sub-buffer.

If a buffer object is constructed from a `cl_mem` object, then the buffer is created and initialized from the OpenCL memory object. The SYCL system may copy the data to another device and/or context, but must copy it back (if modified) at the point of destruction of the buffer. The user must provide a `queue` and `event`. The memory object is assumed to only be available to the SYCL scheduler after the event has signaled and is assumed to be currently resident on the context and device signified by the `queue`.

3.4.3 Images

The class `image<int dimensions>` (Table 3.23) defines shared image data of one, two or three dimensions, that can be used by kernels in queues and has to be accessed using `accessor` classes with image accessor modes.

Image constructors are listed in Table 3.23 and methods in Table 3.24. Where relevant, it is the responsibility of the user to ensure that the format of the data matches the format described by `order` and `type`. Custom image allocators can be defined, but they need to be byte-sized allocators.

If an image object is constructed from a `cl_mem` object, then the image is created and initialized from the OpenCL memory object. The SYCL system may copy the data to the host, but must copy it back (if modified) at the point of destruction of the image. The user must provide a `queue` and `event`. The memory object is assumed to only be available to the SYCL scheduler after the event has signaled and is assumed to be currently resident on the context and device signified by the `queue`.

3.4.3.1 Image Interface

```
1 namespace cl {
2 namespace sycl {
3 namespace image_format {
4 enum class channel_order : unsigned int {
5     R,
6     Rx,
7     A,
8     INTENSITY,
9     LUMINANCE,
10    RG,
11    RGx,
```



```

12  RA,
13  RGB,
14  RGBx,
15  RGBA,
16  ARGB,
17  BGRA
18  }
19
20  enum class channel_type : unsigned int {
21      SNORM_INT8,
22      SNORM_INT16,
23      UNORM_INT8,
24      UNORM_INT16,
25      UNORM_SHORT_565,
26      UNORM_SHORT_555,
27      UNORM_INT_101010,
28      SIGNED_INT8,
29      SIGNED_INT16,
30      SIGNED_INT32,
31      UNSIGNED_INT8,
32      UNSIGNED_INT16,
33      UNSIGNED_INT32,
34      HALF_FLOAT,
35      FLOAT
36  }
37  } // namespace image_format
38
39  template <int dimensions, typename AllocatorT = cl::sycl::image_allocator>
40  class image {
41  public:
42      image(void *hostPointer, image_format::channel_order order,
43            image_format::channel_type type, const range<dimensions> &range);
44
45      image(void *hostPointer, image_format::channel_order order,
46            image_format::channel_type type, const range<dimensions> &range,
47            range<dimensions - 1> &pitch);
48
49      image(shared_ptr_class<void> &hostPointer,
50            image_format::channel_order order, image_format::channel_type type,
51            const range<dimensions> &range, mutex_class *m = nullptr)
52
53      image(shared_ptr_class<void> &hostPointer,
54            image_format::channel_order order, image_format::channel_type type,
55            const range<dimensions> &range, const range<dimensions - 1> &pitch,
56            mutex_class *m = nullptr)
57
58      image(const image<dimensions, AllocatorT> & rhs);
59
60      image<dimensions,AllocatorT> &operator=(
61          const image<dimensions,AllocatorT> &rhs);
62
63      image(cl_mem memObject, const queue &fromQueue, event availableEvent = {});
64
65      image(const image &rhs);
66

```

```

67 ~image();
68
69 range<dimensions> get_range();
70
71 range<dimensions - 1> get_pitch();
72
73 size_t get_size();
74
75 size_t get_count();
76
77 template <access::mode accessMode,
78           access::target accessTarget = access::image>
79 accessor<T, dimensions, accessMode, accessTarget> get_access();
80
81 template<T>
82 set_final_data(weak_ptr_class<T> & finalPointer));
83 };
84 } // namespace sycl
85 } // namespace cl

```

Constructors	Description
<pre> template<int dimensions, typename AllocatorT= cl::sycl::image_allocator> image (void * hostPointer, image_format::channel_order order, image_format::channel_type type, const range<dimensions> & range) </pre>	<p>Create a new image of the given size with storage managed by SYCL.</p>
<pre> template<int dimensions, typename AllocatorT= cl::sycl::image_allocator> image<dimensions>(void *hostPointer, image_format::channel_order order, image_format::channel_type type, const range<dimensions> & range, const range< dimensions-1> &pitch) </pre>	<p>Construct an image. The associated host memory is in <code>hostPointer</code>. The type of the image data is defined by <code>order</code> and <code>type</code>. The size of the image in pixels is defined by <code>size</code> and the pitch is defined by <code>pitch</code>.</p>
<pre> template<int dimensions, typename AllocatorT= cl::sycl::image_allocator> image(shared_ptr_class<void>& hostPointer, image_format::channel_order order, image_format ::channel_type type, const range<dimensions> & range, mutex_class * mutex = nullptr) </pre>	<p>Construct an image from shared host memory between the SYCL application and runtime. The type of the image data is defined by <code>order</code> and <code>type</code>. The size of the image in pixels is defined by <code>size</code>. The mutex is locked by the runtime whenever the data is in use and unlocked otherwise. Data is synchronized with host pointer when the mutex is unlocked by the runtime.</p>

Continued on next page

Table 3.23: Constructors for the `image` class.

Constructors	Description
<pre>template<int dimensions, typename AllocatorT= cl:: sycl::image_allocator> image(shared_ptr_class<void>& hostPointer, image_format::channel_order order, image_format ::channel_type type, const range<dimensions> & range, const range< dimensions-1> & pitch, mutex_class * mutex = nullptr)</pre>	Construct an image using the data in host pointer. The type of the image data is defined by order and type. The size of the image in pixels is defined by size. The pitch of the image data, in bytes, is defined by pitch. The ownership is shared between the SYCL runtime and application. The mutex is locked by the runtime whenever the data is in use and unlocked otherwise. Data is synchronized with hostPointer when the mutex is unlocked by the runtime.
<pre>template<int dimensions, typename AllocatorT= cl:: sycl::image_allocator> image(image<dimensions, AllocatorT> & rhs)</pre>	Copy construct an image as a reference to another image. Images are reference counted, so that they all point to the same underlying memory.
<pre>template<int dimensions, typename AllocatorT= cl:: sycl::image_allocator> image<dimensions>(cl_mem memObject, queue fromQueue, event availableEvent = {})</pre>	Create an image from an existing OpenCL memory object associated with a context after waiting for an event signaling the availability of the OpenCL data. memObject is the OpenCL memory object to use. fromQueue is the queue associated to the memory object. availableEvent specifies the event to wait for if non-null. Retains a reference to the cl_mem object. Caller should release the passed cl_mem object when it is no longer needed. Note that an image created from a cl_mem object will only have one underlying cl_mem for the lifetime of the buffer and use on an incompatible queue constitutes an error.
<pre>image(const image<dimensions, AllocatorT> & rhs)</pre>	Copy Constructor
<pre>image<dimensions,AllocatorT> &operator=(const image<dimensions,AllocatorT> &rhs)</pre>	Assignment operator which will be sharing the same underlying object and will be reference counted internally by the SYCL runtime.
End of table	

Table 3.23: Constructors for the `image` class.

Methods	Description
<pre>range<dimensions> get_range()</pre>	Return a range object representing the size of the image in terms of the number of elements in each dimension as passed to the constructor.
<pre>range<dimensions-1> get_pitch()</pre>	Return a range object representing the pitch of the image in bytes.
Continued on next page	

Table 3.24: Methods of the `image` class.

Methods	Description
<code>size_t get_count()</code>	Returns the total number of elements in the image. Equal to <code>get_range()[0] * ... * get_range()[dimensions-1]</code> .
<code>size_t get_size()</code>	Returns the size of the image storage in bytes. The number of bytes may be greater than <code>get_count()*element size</code> due to padding of elements, rows and slices of the image for efficient access.
<pre>template< access::mode mode, access::target target=access::image> accessor<T, dimensions, mode, target> get_access()</pre>	Returns a valid accessor to the image with the specified access mode and target. The value of target can be <code>access::image</code> or <code>access::host_image</code> .
<pre>template<T> set_final_data(weak_ptr_class<T>& finalPointer))</pre>	Function that sets the final data pointer, to be different than the original pointer given. This is mainly to be used when a <code>shared_ptr</code> is given in the constructor and the output data will reside in a different location than the initialization data.
End of table	

Table 3.24: Methods of the `image` class.

3.4.3.2 Image Synchronization Rules

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 undefined 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 behaviour occurs on image object destruction.

3.4.4 Sharing Host Memory With The SYCL Data Management Classes

In SYCL in order to allow the SYCL runtime to do the 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 at 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. However, the same rules will apply to images, as well.

3.4.4.1 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 in 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 3.4.6. 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 use to copy data back, if needed, before buffer destruction. This host pointer will not de-allocated by the runtime when the buffer is back on the host, and the data is always copied back from the device (if there was a need for it).

3.4.4.2 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 situations 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.

If the pointer contained in the `unique_ptr` is null, pointer is initialized internally in the runtime but no data is copied in. This will be the generic case of a buffer constructor that takes no host pointer.

In this case the buffer and image constructors used will be the following:

- `template<typename T, int dimensions, typename U=default_allocator> buffer(cl::sycl::unique_ptr_class<T> hostData, range<dimensions> myRange);`
- `template<int dimensions, typename U=image_allocator> image(cl::sycl::unique_ptr_class<void> > hostData, cl_channel_order order, cl_channel_type type, range<dimensions> size)`

For example, the following could be used:

```
1 {
2   cl::sycl::unique_ptr_class<int> ptr(data);
3   buffer<int, 1> b(std::move(ptr));
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 `cl::sycl::weak_ptr_class` to enable copying data back, to another host memory address that is going to be valid after buffer construction.

```
1 {
2   cl::sycl::unique_ptr_class<int> ptr(data);
3   buffer<int, 1> b(std::move(ptr));
4   // ptr is not valid anymore
5   // There is nowhere to copy data back
```

```

6   b.set_final_data(cl::sycl::weak_ptr_class<int>(....))
7 }

```

3.4.4.3 Shared SYCL ownership of the host memory

When an `shared_ptr` is passed to the buffer constructor, then the buffer object and the developer's application is shared. If the shared pointer is still used in the application's side then the data are going to be copied back from the buffer or image and be available to the application after the buffer or image is destroyed.

If the data pointed by the shared object is initialized to some data, then that data is used to initialize the buffer. If the shared pointer is null, the pointer is initialized by the runtime internally (and, therefore, the user can use it afterwards in the host).

When the buffer is destroyed, 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 and 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()` method 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()` method has put the data (or not wait nor copy if set final data is `nullptr`).

```

1 {
2   cl::sycl::shared_ptr_class<int> ptr(data);
3   {
4     buffer<int, 1> b(ptr, range<2>(10, 10));
5   } // Data is copied back because there is an user side shared ptr
6 }

```

```

1 {
2   cl::sycl::shared_ptr_class<int> ptr(data);
3   {
4     buffer<int, 1> b(ptr, range<2>(10, 10));
5     ptr.release();
6   } // Data is not copied back, there is no user side shared ptr.
7 }

```

3.4.5 Synchronisation Primitives

When the user wants to use the buffer simultaneously in the SYCL runtime and its own code (e.g. a multi-threaded mechanism) and want to use manual synchronization without host accessors, a pointer to a `cl::sycl::mutex` can be passed to the buffer constructor.

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

```

1 {
2   cl::sycl::mutex_class * m;

```

```

3  auto shD = std::make_shared<int>(42)
4  {
5  buffer<int, 1> b(shD, m);
6
7  m.lock();
8  // User accesses the data
9  do_something(shD);
10 m.unlock();
11
12 }
13 }

```

When the runtime releases the mutex the user is guaranteed to have the data copied back on the shared pointer - unless the final data destination has been changed using the method `set_final_data`.

3.4.6 Accessors

Accessors manage the access to data in buffers and images. The user specifies the type of access to the data and the SYCL implementation ensures that the data is accessible in the right way on the right device in a queue. This separation allows an SYCL implementation to choose an efficient way to provide access to the data within an execution schedule. Common ways of allowing data access to shared data in a heterogeneous system include copying between different memory systems, mapping memory into different device address spaces, or direct sharing of data in memory. The buffers and images are SYCL runtime classes that provide the management of the data.

Accessors are *device accessors* by default, but can optionally be specified as being host accessors. Device accessors can only be constructed within command groups and provide access to the underlying data in a queue. Only a kernel can access data using a device accessor. Constructing a device accessor is a non-blocking operation: the synchronization is added to the queue, not the host.

Host accessors can be created outside command groups and give immediate access to data on the host. Construction of host accessors is blocking, waiting for all previous operations on the underlying buffer or image to complete, including copying from device memory to host memory. Any subsequent device accessors need to block until the processing of the host accessor is done and the data are copied to the device.

Accessors always have an *element data type*. When accessing a buffer, the accessor's element data type must match the same data type as the buffer. An image accessor may have an element data type of either an integer vector or a floating-point vector. The image accessor data type provides the number and type of components of the pixel read. The actual format of the underlying image data is not encoded in the accessor, but in the image object itself. The data types allowed by buffer can be SYCL data types or user defined types that are aligned with the restrictions on kernel parameter passing [3.5.7](#).

There are two enumeration types inside namespace `cl::sycl::access`, `access::mode` and `access::target`. These two enumerations define both the access mode and the data that the accessor is targeting.

3.4.6.1 Access modes

The `mode` enumeration, shown in [Table 3.25](#), has a base value, which must be provided.

The user must provide the *access mode* when defining an accessor. This information is used by the scheduler to ensure that any data dependencies are resolved by enqueueing any data transfers before or after the execution of a kernel. If a command group contains only *discard write mode* accesses to a buffer, then the previous contents of the buffer (or sub-range of the buffer, if provided) are not preserved. 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. A command-group's access to a specific buffer is the union of all access modes to that buffer in the command group, regardless of construction order. Atomic access is only valid to `local`, `global_buffer` and `host_buffer` targets (see next section).

<code>access::mode</code>	Description
<code>read</code>	read-only access
<code>write</code>	write-only access. Previous contents not discarded.
<code>read_write</code>	read and write access
<code>discard_write</code>	write-only access. Previous contents discarded.
<code>discard_read_write</code>	read and write access. Previous contents discarded.
<code>atomic</code>	atomic access.

Table 3.25: Enumeration of access modes available to accessors.

3.4.6.2 Access targets

The target enumeration, shown in Table 3.26, describes the type of object to be accessed via the accessor. The different values of the target enumeration require different constructors for the accessors.

<code>access::target</code>	Description
<code>global_buffer</code>	Access <code>buffer</code> via global memory.
<code>constant_buffer</code>	Access <code>buffer</code> via constant memory.
<code>local</code>	Access work-group-local memory.
<code>image</code>	Access an <code>image</code> .
<code>host_buffer</code>	Access a <code>buffer</code> immediately in host code.
<code>host_image</code>	Access an <code>image</code> immediately in host code.
<code>image_array</code>	Access an array of <code>images</code> on a device.

Table 3.26: Enumeration of access modes available to accessors.

3.4.6.3 Accessor class

The accessor makes a data available to host code, or to a specific kernel. The `accessor` is parameterized with the type and number of dimensions of the data. An accessor also has a mode, which defines the operations possible on the underlying data (see Table 3.25) and a target (see Table 3.26, which defines the type of data object to be modified. The constructors and methods available on an accessor depend on the mode and target.

The generic methods for the `accessor` class are defined in Table 3.28. Available methods are limited by the access mode and target provided as template parameters to the accessor object.

```
1 namespace cl {
2 namespace sycl {
3 namespace access {
4 enum class mode {
5     read = 1,
6     write,
7     read_write,
8     discard_write,
9     discard_read_write,
10    atomic
11 };
12
13 enum class target {
14     global_buffer = 2014,
15     constant_buffer,
16     local,
17     image,
18     host_buffer,
19     host_image,
20     image_array
21 };
22 } // namespace access
23
24 template <
25     typename elementType,
26     int dimensions,
27     access::mode accessMode,
28     access::target accessTarget>
29 class accessor {
30 public:
31     using value_type = T;
32     using reference = value_type&;
33     using const_reference = const value_type&;
34
35     // Available only for: access::global_buffer, access::host_buffer,
36     // and access::constant_buffer
37     accessor(buffer<elementType, dimensions> &bufferRef,
38             handler &commandGroupHandler);
39
40     accessor(buffer<elementType, dimensions> &bufferRef,
41             handler &commandGroupHandler,
42             range<dimensions> offset, range<dimensions> range);
43
44     // Available only for: access::image and access::host_image
45     accessor(image<dimensions> &imageRef, handler &commandGroupHandler);
46
47     // Available only for: access::local
48     accessor(range<dimensions> allocationSize, handler &commandGroupHandler);
49
50     size_t get_size();
51
52     // Methods available to buffer targets
```

```

53 // Available when access_mode includes non-atomic write permissions
54 // and dimensions>0
55 elementType &operator[](id<dimensions>);
56 // Available when access_mode is read-only and dimensions>0
57 const elementType &operator[](id<dimensions>);
58 // Available when access_mode includes non-atomic write permissions
59 // and dimensions==0
60 elementType &operator*();
61
62 // Available when access_mode is read-only and dimensions==0
63 const elementType &operator[]();
64 elementType operator*();
65
66 // Available when dimensions==0 and access mode is non-atomic.
67 operator elementType();
68
69 // Methods available for image targets
70 __undefined__ &operator()(sampler sample);
71 __undefined__ &operator[](id<dimensions>);
72
73 // Available when the accessor is to an image array
74 // Returns an accessor to a particular slice
75 accessor<elementType, 2, mode, image> operator[](size_t index)
76
77 // Available when the access target is global_buffer, host_buffer,
78 // or local, and access mode is atomic
79 // and dimensions>0
80 atomic<elementType> &operator[](id<dimensions>);
81
82 // Available when the access target is global_buffer, host_buffer,
83 // or local, and access mode is atomic
84 // and dimensions==0
85 atomic<elementType> operator()();
86 atomic<elementType> operator*();
87
88
89 //Available when access target is local and access mode is non-atomic
90 local_ptr<elementType> get_pointer();
91
92 //Available when access target is global_buffer and access mode is non-atomic
93 global_ptr<elementType> get_pointer();
94
95 //Available when access target is constant_buffer
96 constant_ptr<elementType> get_pointer();
97
98 //Available when access target is host_buffer and access mode is non-atomic
99 elementType * get_pointer();
100 };
101
102 } // namespace sycl
103 } // namespace cl

```

Constructors	Description
<pre>accessor(buffer<elementType, dimensions> &bufferRef, handler &commandGroupHandler)</pre>	<p>Construct a buffer accessor from a buffer using a command group handler object from the command group scope. Constructor only available for access modes <code>global_buffer</code>, <code>host_buffer</code> <code>constant_buffer</code> see Table 3.25. <code>access_target</code> defines the form of access being obtained. See Table 3.26.</p>
<pre>accessor(buffer<elementType, dimensions> &bufferRef, handler &commandGroupHandler, range<dimensions> offset, range<dimensions> range)</pre>	<p>Construct a buffer accessor from a buffer given a specific range for access permissions and an offset that provides the starting point for the access range using a command group handler object from the command group scope. This accessor limits the processing of the buffer to the <code>[offset, offset+range]</code> for every dimension. Any other parts of the buffer will be unaffected. Constructor only available for access modes <code>global_buffer</code>, <code>host_buffer</code> or <code>constant_buffer</code> (see Table 3.25). <code>access_target</code> defines the form of access being obtained (see Table 3.26). This accessor is recommended for <code>discard_write</code> and <code>discard_read_write</code> access modes, when the unaffected parts of the processing should be retained.</p>
<pre>accessor(image<dimensions> &imageRef, handler &commandGroupHandler)</pre>	<p>Construct an image accessor from an image using a command group handler object from the command group scope. Constructor only available if <code>accessMode</code> is <code>image</code>, or <code>host_image</code>, see Table 3.25. <code>access_target</code> defines the form of access being obtained. See Table 3.26. The <code>elementType</code> for image accessors must be defined by the user and is the type returned by any sampler or accessor read operation, as well as the value accepted by any write operation. It must be an <code>int</code>, <code>unsigned int</code> or <code>float</code> vector of 4 dimensions.</p>
<pre>accessor(range<dimensions> allocationSize, handler &commandGroupHandler)</pre>	<p>Construct an accessor of dimensions dimensions with elements of type <code>elementType</code> using the passed range to specify the size in each dimension. It needs as a parameter a command group handler object from the command group scope. Constructor only available if <code>accessMode</code> is <code>local</code>, see Table 3.25.</p>
End of table	

Table 3.27: Accessor constructors.

Methods	Description
<code>size_t get_size()</code>	Returns the size of the underlying buffer in number of elements.
<code>elementType &operator[](id<dimensions>)</code>	Return a writeable reference to an element in the buffer. Available when mode includes non-atomic write permissions and <code>dimensions == 0</code> .
<code>elementType &operator[](int)</code>	Return a writeable reference to an element in the buffer. Available when mode includes non-atomic write permissions and <code>dimensions == 1</code> .
<code>const elementType &operator[](id<dimensions>)</code>	Return the value of an element in the buffer. Available when mode is read-only and <code>dimensions == 0</code> .
<code>const elementType &operator[](int)</code>	Return the value of an element in the buffer. Available when mode is read-only and <code>dimensions == 1</code> .
<code>elementType &operator[]()</code>	Return a writeable reference to the element in the buffer. Available when mode includes non-atomic write permissions and <code>dimensions == 0</code> .
<code>const elementType &operator[]()</code>	Return the value of the element in the buffer. Available when mode is read-only and <code>dimensions == 0</code> .
<code>operator elementType()</code>	Return the value of the element in the buffer. Available when mode is non-atomic and <code>dimensions == 0</code> .
<code>accessor<elementType, 2, mode, image> operator[](size_t index)</code>	Returns an accessor to a particular plane of an image array. Available when accessor acts on an image array.
<code>__undefined__ <dimensions-1> &operator[](int)</code>	Return an intermediate type with an additional subscript operator for each subsequent dimension of buffer where (<code>dimensions > 0</code>). Available when mode non-atomic and for access mode read only the return type is <code>const</code> .
<code>__undefined__ &operator()(sampler sample)</code>	Return the value of an element in the image given a sampler. Available only for the case of an image accessor type.
<code>__undefined__ &operator [] (id < dimensions >)</code>	Return the value of an element in the image with a sampler-less read. Available only for the case of an image accessor type.
<code>atomic<elementType> &operator[](id<dimensions>)</code>	Returns a reference to an atomic object, when the accessor is of type <code>access::global_buffer</code> , <code>access::local_buffer</code> , <code>access::host_buffer</code> , the target mode is <code>access::mode::atomic</code> and <code>dimensions > 0</code> .

Continued on next page

Table 3.28: Methods for the `accessor` class.

Methods	Description
<code>atomic<elementType> &operator()()</code>	Returns a reference to an atomic object, when the accessor is of type <code>access::global_buffer</code> , <code>access::local_buffer</code> , <code>access::host_buffer</code> , the target mode is <code>access::mode::atomic</code> and <code>dimensions==0</code> .
<code>atomic<elementType> &operator*()</code>	Returns a reference to an atomic object, when the accessor is of type <code>access::global_buffer</code> , <code>access::local_buffer</code> , <code>access::host_buffer</code> , the target mode is <code>access::mode::atomic</code> and <code>dimensions==0</code> .
<code>local_ptr<elementType> get_pointer()</code>	Returns the accessor pointer, when the accessor is of type <code>access::local</code> and mode is non-atomic.
<code>global_ptr<elementType> get_pointer()</code>	Returns the accessor pointer, when the accessor is of type <code>access::global_buffer</code> and mode is non-atomic.
<code>constant_ptr<elementType> get_pointer()</code>	Returns the accessor pointer, when the accessor is of type <code>access::constant_buffer</code>
<code>elementType* get_pointer()</code>	Returns the accessor pointer, when the accessor is of type <code>access::host_buffer</code> and mode is non-atomic.
End of table	

Table 3.28: Methods for the `accessor` class.

3.4.6.4 Buffer accessors

Accessors to buffers are constructed from a buffer with the same element data type and dimensionality as the accessor. A buffer accessor uses *global* memory by default, but can optionally be set to use *constant* memory. Accessors that use constant memory are restricted by the underlying OpenCL restrictions on device constant memory, i.e. there is a maximum total constant memory usable by a kernel and that maximum is specified by the OpenCL device. Only certain methods and constructors are available for buffer accessors.

The array operator `[id<dimensions>]` provides access to the elements of the buffer. The user can provide an index as an `id` parameter of the same dimensionality of the buffer, or just like C++ arrays, can provide one array operator per dimension, with individual indices of type `size_t` (e.g. `myAccessor[i][j][k]`).

The address space for the index operator matches that of the accessor target. For an `access::global_buffer`, the address space is `global`. For an `access::constant_buffer`, the address space is `global`.

Accessors to buffers can be constructed to only access a *sub-range* of the buffer. The sub-range restricts access to just that range of the buffer, which the scheduler can use as extra information to extract more parallelism from queues as well as restrict the amount of information copied between devices.

3.4.6.5 Image accessors

Accessors that target images must be constructed from images of the same dimensionality as the accessor. The `target` parameter must be either `image` or `host_image`. The `dataType` parameter must be a 4 dimension vector of `unsigned int`, `int` or `float`. The array operator `[id<dimensions>]` provides samplerless reading and writing of the image. The user can provide an index as an `id` parameter of the same dimensionality of the image, or just like C++ arrays, can provide one array operator per dimension, with individual indices of type `size_t` (e.g. `myAccessor[i][j][k]`). The bracket operator takes a `sampler` (see 3.4.8) parameter, which then allows floating-point sampler-based reading using the array operator (e.g. `myAccessor(mySampler)[my2dFloatVector]`).

To enable the reading and writing of pixels with and without samplers, using standard C++ operators, there are two internal classes: `__image_ref` and `__sampler`. These classes only exist to ensure that assignment to pixels uses image write functions and reading the value of pixels uses image read functions.

There are restrictions that apply to `cl::sycl::access::mode` depending on the SYCL device they are used. In OpenCL 1.2 devices there are extensions that can allow `access::mode::read_write` on images, but this is not a core feature of the OpenCL 1.2 devices. It is the developer's responsibility to check whether there is this capability on the target SYCL device. Core capabilities are `access::mode::read`, `access::mode::write` and `access::discard_write`.

3.4.6.6 Local accessors

Accessors can also be created for *local* memory, to enable pre-allocation of local buffers used inside a kernel. These accessors are constructed using `cl::sycl::range`, which defines the size of the memory to be allocated on a per work-group basis and must be constructed with an access target of `local`. Local memory is only shared across a work-group. A local accessor can provide a `local_ptr` to the underlying data within a kernel and is only usable within a kernel. The host has no access to the data of the local buffer and cannot read or write to the data, so the accessor cannot read or write data back to the host. There can be no associated host pointer for a local buffer or data transfers.

Local accessors are not valid for single-task or basic `parallel_for` invocations.

3.4.6.7 Host accessors

Host accessors have a target of `access::host_buffer` or `access::host_image`. Unlike other accessors, host accessors should be constructed outside of any command group and they do not require an object of a command group `handler`. The constructor will block until the data is ready for the host to access, while the destructor will block any further operations on the data in any SYCL queue. There are no special constructor or method signatures for host accessors, so there are no table for special host accessors here (see buffer and image accessors above).

Host accessors are constructed outside command groups and not associated with any queue, so any error reporting is synchronous. By default, error reporting is via exceptions.

3.4.6.8 Accessor capabilities and restrictions

Accessors provide access on the device or on the host to a buffer or image. The access modes allowed depend on the accessor type and target. A device accessor grants access to a kernel inside a command group scope, and depending on the access target, there are different accesses allowed. A host accessor grants access to the host program to the access target. Tables 3.29, 3.30 and 3.31 show all the permitted access modes depending on target.

Accessor Type	Access Target	Access mode	Data Type	Description
Device	global_buffer	read write read_write discard_write discard_read_write	All available data types supported in SYCL.	Access a buffer allocated in global memory on the device.
Device	constant_buffer	read	All available data types supported in SYCL.	Access a buffer allocated in constant memory on the device.
Host	host_buffer	read write read_write discard_write discard_read_write	All available data types supported in SYCL.	Access a host allocated buffer on host.
Device	local	read write read_write	All supported data types in local memory	Access work-group local buffer, which is not associated with a host buffer. This is only accessible on device.

Table 3.29: Description of all the accessor types and modes with their valid combinations for buffers and local memory

Accessor Type	Access Target	Access mode	Data Type	Description
Device	image	read write read_write discard_ writediscard_read_ write	uint4, int4, float4, half4	Access an image on device.
Host	host_image	read write read_write discard_ writediscard_read_ write	uint4, int4, float4, half4	Access an image on the host.
Device	image_array	read write read_write discard_ writediscard_read_ write	uint4, int4, float4, half4	Access an array of images on device.

Table 3.30: Description of all the accessor types and modes with their valid combinations for images

Rules for casting apply to the accessors, as there is only a specific set of permitted conversions.

Accessor Types	Original Accessor Target	Original Access Mode	Converted Accessor Target	Converted Access Mode
Device	global_buffer	read_write	global_buffer	read write discard_read_write
Device	local_buffer	read_write	local_buffer	read write
Host	host_buffer	read_write	host_buffer	read write discard_read_write

Table 3.31: Description of the accessor to accessor conversions allowed

3.4.7 Address space classes

In OpenCL, there are four different address spaces. These are: global, local, constant and private. In OpenCL C, these address spaces are manually specified using OpenCL-specific keywords. In SYCL, the device compiler is expected to auto-deduce the address space for pointers in common situations of pointer usage. However, there are situations where auto-deduction is not possible. Here are the most common situations:

- When linking SYCL kernels with OpenCL C functions. In this case, it is necessary to specify the address space for any pointer parameters when declaring an `extern "C"` function.
- When declaring data structures with pointers inside, it is not possible for the SYCL compiler to deduce at the time of declaration of the data structure what address space pointer values assigned to members of the structure will be. So, in this case, the address spaces will have to be explicitly declared by the developer.
- When a pointer is declared as a variable, but not initialized, then address space deduction is not automatic and so an explicit pointer class should be used, or the pointer should be initialized at declaration.

3.4.7.1 Explicit pointer classes

Explicit pointer classes are just like pointers: they can be converted to and from pointers with compatible address spaces, qualifiers and types. Assignment between explicit pointer types of incompatible address spaces is illegal. In SYCL 1.2, all address spaces are incompatible with all other address spaces. For a future SYCL 2.0, a generic address space will be compatible with all other address spaces. Conversion from an explicit pointer to a C++ pointer preserves the address space.

In order to facilitate SYCL/OpenCL C interoperability, the `pointer_t` type is provided. It is an implementation defined type which corresponds to the underlying OpenCL C pointer type and can be used in `extern "C"` function declarations for OpenCL functions used in SYCL kernels.

Explicit Pointer Classes	OpenCL Address Space	Compatible Accessor Target
<code>global_ptr</code>	<code>__global</code>	<code>global_buffer</code>
<code>constant_ptr</code>	<code>__constant</code>	<code>constant_buffer</code>
<code>local_ptr</code>	<code>__local</code>	<code>local</code>
<code>private_ptr</code>	<code>__private</code>	<code>none</code>
End of table		

Table 3.32: Description of the pointer classes

An overview of the interface provided for all the explicit pointer classes is the following, for the full interface please refer to [B.1](#).

```
1 namespace cl {
2 namespace sycl {
3
4 template <typename ElementType>
5 class global_ptr {
6 public:
7     // Implementation defined pointer type that corresponds to the SYCL/OpenCL
8     // interoperability type for OpenCL C functions
```

```

9   typedef __undefined__ pointer_t;
10
11  global_ptr(pointer_t); // global pointer
12
13  template <access::mode Mode>
14  global_ptr(accessor<ElementType, 1, Mode, global_buffer>);
15
16  global_ptr(const global_ptr&);
17
18  ~global_ptr();
19
20  ElementType& operator*();
21  ElementType& operator[](size_t i);
22
23  // implementation defined implicit conversion
24  // to OpenCL C pointer types.
25  operator pointer_t();
26
27 };
28
29 template <typename ElementType>
30 class constant_ptr {
31 public:
32     // Implementation defined pointer type that corresponds to the SYCL/OpenCL
33     // interoperability type for OpenCL C functions
34     typedef __undefined__ pointer_t;
35
36     constant_ptr(pointer_t); // constant pointer
37
38     template <access::mode Mode>
39     constant_ptr(accessor<ElementType, 1, Mode, constant_buffer>);
40
41     ElementType& operator*();
42     ElementType& operator[](size_t i);
43
44     // implementation defined implicit conversion
45     // to OpenCL C pointer types.
46     operator pointer_t();
47 };
48
49 template <typename ElementType>
50 class local_ptr {
51 public:
52     // Implementation defined pointer type that corresponds to the SYCL/OpenCL
53     // interoperability type for OpenCL C functions
54     typedef __undefined__ pointer_t;
55
56     local_ptr(pointer_t); // local pointer
57
58     template <access::mode Mode>
59     local_ptr(accessor<ElementType, 1, Mode, local_buffer>);
60
61     ElementType& operator*();
62     ElementType& operator[](size_t i);
63

```

```

64 // implementation defined implicit conversion
65 // to OpenCL C pointer types.
66 operator pointer_t();
67 };
68
69 template <typename ElementType>
70 class private_ptr {
71 public:
72 // Implementation defined pointer type that corresponds to the SYCL/OpenCL
73 // interoperability type for OpenCL C functions
74 typedef __undefined__ pointer_t;
75
76 private_ptr(pointer_t); // private pointer
77
78 ElementType& operator*();
79 ElementType& operator[](size_t i);
80
81 // implementation defined implicit conversion
82 // to OpenCL C pointer types.
83 operator pointer_t();
84 };
85 } // namespace sycl
86 } // namespace cl

```

Constructors	Description
template <typename ElementType> global_ptr(pointer_t)	Constructs a global_ptr from the underlying ElementType pointer.
template <access::mode Mode> global_ptr(accessor<ElementType, 1, Mode, global_buffer>)	Constructs a global_ptr from an accessor of <code>access::target::global_buffer</code> .
template <typename ElementType> global_ptr(const global_ptr &)	Copy constructor.
End of table	

Table 3.33: Constructors for [global_ptr](#) explicit pointer class.

Operators	Description
template <typename ElementType> ElementType &operator*()	Returns the ElementType of the de-referenced pointer class.
template <typename ElementType> ElementType &operator[](size_t i)	Subscript operator.
template <typename ElementType> operator pointer_t()	Returns the underlying pointer type of the global_ptr class
End of table	

Table 3.34: Operators on the [global_ptr](#) explicit pointer class.

Constructors	Description
<code>template <typename ElementType> constant_ptr(pointer_t)</code>	Constructs a <code>constant_ptr</code> from the underlying <code>ElementType</code> pointer.
<code>template <access::mode Mode> constant_ptr(accessor<ElementType, 1, Mode, constant_buffer>)</code>	Constructs a <code>constant_ptr</code> from an accessor of <code>access::target::constant_buffer</code> .
<code>template <typename ElementType> constant_ptr(const constant_ptr &)</code>	Copy constructor.
End of table	

Table 3.35: Constructors for `constant_ptr` explicit pointer class.

Operators	Description
<code>template <typename ElementType> ElementType &operator*()</code>	Returns the <code>ElementType</code> of the dereferenced pointer class.
<code>template <typename ElementType> ElementType &operator[](size_t i)</code>	Subscript operator.
<code>template <typename ElementType> operator pointer_t()</code>	Returns the underlying pointer type of the <code>constant_ptr</code> class
End of table	

Table 3.36: Operators on the `constant_ptr` explicit pointer class.

Constructors	Description
<code>template <typename ElementType> local_ptr(pointer_t)</code>	Constructs a <code>local_ptr</code> from the underlying <code>ElementType</code> pointer.
<code>template <access::mode Mode> local_ptr(accessor<ElementType, 1, Mode, constant_buffer>)</code>	Constructs a <code>local_ptr</code> from an accessor of <code>access::target::local</code> .
<code>template <typename ElementType> local_ptr(const local_ptr &)</code>	Copy constructor.
End of table	

Table 3.37: Constructors for `local_ptr` explicit pointer class.

Operators	Description
<code>template <typename ElementType> ElementType &operator*()</code>	Returns the <code>ElementType</code> of the dereferenced pointer class.
<code>template <typename ElementType> ElementType &operator[](size_t i)</code>	Subscript operator.
<code>template <typename ElementType> operator pointer_t()</code>	Returns the underlying pointer type of the <code>local_ptr</code> class
End of table	

Table 3.38: Operators on the `local_ptr` explicit pointer class.

Constructors	Description
<code>template <typename ElementType> private_ptr(pointer_t)</code>	Constructs a <code>private_ptr</code> from an ElementType pointer.
<code>template <typename ElementType> private_ptr(const private_ptr &)</code>	Copy constructor.
End of table	

Table 3.39: Constructors for `private_ptr` explicit pointer class.

Operators	Description
<code>template <typename ElementType> ElementType &operator*()</code>	Returns the ElementType of the dereferenced pointer class.
<code>template <typename ElementType> ElementType &operator[](size_t i)</code>	Subscript operator.
<code>template <typename ElementType> operator pointer_t()</code>	Returns the underlying pointer type of the <code>private_ptr</code> class
End of table	

Table 3.40: Operators on the `private_ptr` explicit pointer class.

Non-member functions	Description
<code>template <typename ElementType> bool operator==(const global_ptr<ElementType>& lhs, const global_ptr<ElementType>& rhs)</code>	Comparison operator <code>==</code> for <code>global_ptr</code> class.
<code>template <typename ElementType> bool operator!=(const global_ptr<ElementType>& lhs, const global_ptr<ElementType>& rhs)</code>	Comparison operator <code>!=</code> for <code>global_ptr</code> class.
<code>template <typename ElementType> bool operator<(const global_ptr<ElementType>& lhs, const global_ptr<ElementType>& rhs)</code>	Comparison operator <code><</code> for <code>global_ptr</code> class.
<code>template <typename ElementType> bool operator>(const global_ptr<ElementType>& lhs, const global_ptr<ElementType>& rhs)</code>	Comparison operator <code>></code> for <code>global_ptr</code> class.
<code>template <typename ElementType> bool operator>=(const global_ptr<ElementType>& lhs, const global_ptr<ElementType>& rhs)</code>	Comparison operator <code>>=</code> for <code>global_ptr</code> class.
<code>template <typename ElementType> bool operator<=(const global_ptr<ElementType>& lhs, const global_ptr<ElementType>& rhs)</code>	Comparison operator <code><=</code> for <code>global_ptr</code> class.
Continued on next page	

Table 3.41: Non-member functions of the explicit pointer classes.

Non-member functions	Description
<pre>template <typename ElementType> bool operator!=(const global_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator != for <code>global_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator!=(nullptr_t lhs, const global_ptr< ElementType>& rhs)</pre>	Comparison operator != for <code>global_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator==(const global_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator == for <code>global_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator==(nullptr_t lhs, const global_ptr< ElementType>& rhs)</pre>	Comparison operator == for <code>global_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>(const global_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator > for <code>global_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>(nullptr_t lhs, const global_ptr< ElementType>& rhs)</pre>	Comparison operator > for <code>global_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<(const global_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator < for <code>global_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<(nullptr_t lhs, const global_ptr< ElementType>& rhs)</pre>	Comparison operator < for <code>global_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator <=(const global_ptr < ElementType > &lhs, nullptr_t rhs)</pre>	Comparison operator >= for <code>global_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>=(nullptr_t lhs, const global_ptr< ElementType>& rhs)</pre>	Comparison operator >= for <code>global_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<=(const global_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator <= for <code>global_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<=(nullptr_t lhs, const global_ptr< ElementType>& rhs)</pre>	Comparison operator <= for <code>global_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator==(const constant_ptr<ElementType>& lhs, const constant_ptr<ElementType>& rhs)</pre>	Comparison operator == for <code>constant_ptr</code> class.
<pre>template <typename ElementType> bool operator!=(const constant_ptr<ElementType>& lhs, const constant_ptr<ElementType>& rhs)</pre>	Comparison operator != for <code>constant_ptr</code> class.

Continued on next page

Table 3.41: Non-member functions of the explicit pointer classes.

Non-member functions	Description
<pre>template <typename ElementType> bool operator<(const constant_ptr<ElementType>& lhs, const constant_ptr<ElementType>& rhs)</pre>	Comparison operator < for <code>constant_ptr</code> class.
<pre>template <typename ElementType> bool operator>(const constant_ptr<ElementType>& lhs, const constant_ptr<ElementType>& rhs)</pre>	Comparison operator > for <code>constant_ptr</code> class.
<pre>template <typename ElementType> bool operator>=(const constant_ptr<ElementType>& lhs, const constant_ptr<ElementType>& rhs)</pre>	Comparison operator >= for <code>constant_ptr</code> class.
<pre>template <typename ElementType> bool operator<=(const constant_ptr<ElementType>& lhs, const constant_ptr<ElementType>& rhs)</pre>	Comparison operator <= for <code>constant_ptr</code> class.
<pre>template <typename ElementType> bool operator!=(const constant_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator <= for <code>constant_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator!=(nullptr_t lhs, const constant_ptr<ElementType>& rhs)</pre>	Comparison operator != for <code>constant_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator==(const constant_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator == for <code>constant_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator==(nullptr_t lhs, const constant_ptr<ElementType>& rhs)</pre>	Comparison operator == for <code>constant_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>(const constant_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator > for <code>constant_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>(nullptr_t lhs, const constant_ptr <ElementType>& rhs)</pre>	Comparison operator > for <code>constant_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<(const constant_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator < for <code>constant_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<(nullptr_t lhs, const constant_ptr <ElementType>& rhs)</pre>	Comparison operator < for <code>constant_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>=(const constant_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator >= for <code>constant_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>=(nullptr_t lhs, const constant_ptr<ElementType>& rhs)</pre>	Comparison operator >= for <code>constant_ptr</code> class with a <code>nullptr_t</code> .

Continued on next page

Table 3.41: Non-member functions of the explicit pointer classes.

Non-member functions	Description
<pre>template <typename ElementType> bool operator<=(const constant_ptr<ElementType>&& lhs, nullptr_t rhs)</pre>	Comparison operator <= for <code>constant_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<=(nullptr_t lhs, const constant_ptr<ElementType>& rhs)</pre>	Comparison operator <= for <code>constant_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator==(const local_ptr<ElementType>& lhs, const local_ptr<ElementType>& rhs)</pre>	Comparison operator == for <code>local_ptr</code> class.
<pre>template <typename ElementType> bool operator!=(const local_ptr<ElementType>& lhs, const local_ptr<ElementType>& rhs)</pre>	Comparison operator != for <code>local_ptr</code> class.
<pre>template <typename ElementType> bool operator<(const local_ptr<ElementType>& lhs , const local_ptr<ElementType>& rhs)</pre>	Comparison operator < for <code>local_ptr</code> class.
<pre>template <typename ElementType> bool operator>(const local_ptr<ElementType>& lhs , const local_ptr<ElementType>& rhs)</pre>	Comparison operator > for <code>local_ptr</code> class.
<pre>template <typename ElementType> bool operator>=(const local_ptr<ElementType>& lhs, const local_ptr<ElementType>& rhs)</pre>	Comparison operator >= for <code>local_ptr</code> class.
<pre>template <typename ElementType> bool operator<=(const local_ptr<ElementType>& lhs, const local_ptr<ElementType>& rhs)</pre>	Comparison operator <= for <code>local_ptr</code> class.
<pre>template <typename ElementType> bool operator!=(const local_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator != for <code>local_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator!=(nullptr_t lhs, const local_ptr< ElementType>& rhs)</pre>	Comparison operator != for <code>local_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator==(const local_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator == for <code>local_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator==(nullptr_t lhs, const local_ptr< ElementType>& rhs)</pre>	Comparison operator == for <code>local_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>(const local_ptr<ElementType>& lhs , nullptr_t rhs)</pre>	Comparison operator > for <code>local_ptr</code> class with a <code>nullptr_t</code> .

Continued on next page

Table 3.41: Non-member functions of the explicit pointer classes.

Non-member functions	Description
<pre>template <typename ElementType> bool operator>(nullptr_t lhs, const local_ptr< ElementType>& rhs)</pre>	Comparison operator > for <code>local_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<(const local_ptr<ElementType>& lhs , nullptr_t rhs)</pre>	Comparison operator < for <code>local_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<(nullptr_t lhs, const local_ptr< ElementType>& rhs)</pre>	Comparison operator < for <code>local_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>=(const local_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator >= for <code>local_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>=(nullptr_t lhs, const local_ptr< ElementType>& rhs)</pre>	Comparison operator >= for <code>local_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<=(const local_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator <= for <code>local_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<=(nullptr_t lhs, const local_ptr< ElementType>& rhs)</pre>	Comparison operator <= for <code>local_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator==(const private_ptr<ElementType>& lhs, const private_ptr<ElementType>& rhs)</pre>	Comparison operator == for <code>private_ptr</code> class.
<pre>template <typename ElementType> bool operator!=(const private_ptr<ElementType>& lhs, const private_ptr<ElementType>& rhs)</pre>	Comparison operator != for <code>private_ptr</code> class.
<pre>template <typename ElementType> bool operator<(const private_ptr<ElementType>& lhs, const private_ptr<ElementType>& rhs)</pre>	Comparison operator < for <code>private_ptr</code> class.
<pre>template <typename ElementType> bool operator>(const private_ptr<ElementType>& lhs, const private_ptr<ElementType>& rhs)</pre>	Comparison operator > for <code>private_ptr</code> class.
<pre>template <typename ElementType> bool operator>=(const private_ptr<ElementType>& lhs, const private_ptr<ElementType>& rhs)</pre>	Comparison operator >= for <code>private_ptr</code> class.
<pre>template <typename ElementType> bool operator<=(const private_ptr<ElementType>& lhs, const private_ptr<ElementType>& rhs)</pre>	Comparison operator <= for <code>private_ptr</code> class.

Continued on next page

Table 3.41: Non-member functions of the explicit pointer classes.

Non-member functions	Description
<pre>template <typename ElementType> bool operator!=(const private_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator != for <code>local_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator!=(nullptr_t lhs, const private_ptr <ElementType>& rhs)</pre>	Comparison operator != for <code>private_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator==(const private_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator == for <code>private_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator==(nullptr_t lhs, const private_ptr <ElementType>& rhs)</pre>	Comparison operator == for <code>private_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>(const private_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator > for <code>private_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>(nullptr_t lhs, const private_ptr< ElementType>& rhs)</pre>	Comparison operator > for <code>private_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<(const private_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator < for <code>private_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<(nullptr_t lhs, const private_ptr< ElementType>& rhs)</pre>	Comparison operator < for <code>private_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>=(const private_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator >= for <code>private_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator>=(nullptr_t lhs, const private_ptr <ElementType>& rhs)</pre>	Comparison operator >= for <code>private_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<=(const private_ptr<ElementType>& lhs, nullptr_t rhs)</pre>	Comparison operator <= for <code>private_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType> bool operator<=(nullptr_t lhs, const private_ptr <ElementType>& rhs)</pre>	Comparison operator <= for <code>private_ptr</code> class with a <code>nullptr_t</code> .
End of table	

Table 3.41: Non-member functions of the explicit pointer classes.

3.4.7.2 Multi-pointer class

There are situations where a user may want to template a datastructure by an address space. Or, a user may want to write templates that adapt to the address space of a pointer. 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 order to facilitate SYCL/OpenCL C interoperability, the `pointer_t` type is provided. It is an implementation defined type which corresponds to the underlying OpenCL C pointer type and can be used in extern ‘C’ function declarations for OpenCL

functions used in SYCL kernels.

An overview of the interface provided for the `multi_ptr` class is the following, for the full interface please refer to [B.2](#).

```
1 namespace cl {
2 namespace sycl {
3 namespace access {
4 enum class address_space : int {
5     global_space,
6     local_space,
7     constant_space,
8     private_space
9 };
10 } // namespace access
11
12 template <typename ElementType, access::address_space Space>
13 class multi_ptr {
14 public :
15     // Implementation defined pointer type that corresponds to the SYCL/OpenCL
16     // interoperability type for OpenCL C functions
17     typedef __undefined__ pointer_t;
18
19     const address_space space;
20
21     multi_ptr(pointer_t);
22     multi_ptr(const multi_ptr&);
23     ~multi_ptr();
24
25     ElementType &operator*();
26     ElementType &operator[](size_t i);
27
28     // Only if Space == global_space
29     operator global_ptr<ElementType>();
30     global_ptr<ElementType> pointer();
31
32     // Only if Space == local_space
33     operator local_ptr<ElementType>();
34     local_ptr<ElementType> pointer();
35
36     // Only if Space == constant_space
37     operator constant_ptr<ElementType>();
38     constant_ptr<ElementType> pointer();
39
40     // Only if Space == private_space
41     operator private_ptr<ElementType>();
42     private_ptr<ElementType> pointer();
43 };
44
45 template <typename ElementType, access::address_space Space>
46 multi_ptr<ElementType, Space> make_ptr(pointer_t);
47 } // namespace sycl
48 } // namespace cl
```

Constructors	Description
<pre>template <typename ElementType, enum address_space Space> explicit multi_ptr(pointer_t)</pre>	Constructor that takes as an argument a pointer of ElementType.
<pre>template <typename ElementType, access::address_space Space> multi_ptr(const multi_ptr &)</pre>	Copy constructor.
<pre>template <typename ElementType, access::address_space Space> multi_ptr<ElementType, Space> make_ptr(pointer_t))</pre>	Global function to create a <code>multi_ptr</code> instance depending on the address space of the <code>pointer_t</code> pointer type.
End of table	

Table 3.42: Constructors for `multi_ptr` class

Methods	Description
<pre>template <typename ElementType, access::address_space Space> ElementType &operator*()</pre>	Operator that returns a reference to the ElementType of the <code>multi_ptr</code> class.
<pre>template <typename ElementType, access::address_space Space> ElementType &operator[](size_t i)</pre>	Subscript operator.
<pre>template <typename ElementType, access::address_space Space = access::address_space::global_space> operator global_ptr<ElementType>()</pre>	Conversion operator from <code>multi_ptr</code> <ElementType, address_space::global_space> to <code>global_ptr</code> <ElementType>.
<pre>template <typename ElementType, access::address_space Space = access::address_space::global_space> global_ptr<ElementType> pointer()</pre>	Pointer method that returns a <code>global_ptr</code> <ElementType> from a <code>multi_ptr</code> <ElementType, address_space::global_space>.
<pre>template <typename ElementType, access::address_space Space = access::address_space::local_space> operator local_ptr<ElementType>()</pre>	Conversion operator from <code>multi_ptr</code> <ElementType, address_space::local_space> to <code>local_ptr</code> <ElementType>.
<pre>template <typename ElementType, access::address_space Space = access::address_space::local_space> local_ptr<ElementType> pointer()</pre>	Pointer method that returns a <code>local_ptr</code> <ElementType> from a <code>multi_ptr</code> <ElementType, address_space::local_space>.
<pre>template <typename ElementType, access::address_space Space = access::address_space::constant_space> operator constant_ptr<ElementType>()</pre>	Conversion operator from <code>multi_ptr</code> <ElementType, address_space::constant_space> to <code>constant_ptr</code> <ElementType>.
<pre>template <typename ElementType, access::address_space Space = access::address_space::constant_space> constant_ptr<ElementType> pointer()</pre>	Pointer method that returns a <code>constant_ptr</code> <ElementType> from a <code>multi_ptr</code> <ElementType, address_space::constant_space>.
Continued on next page	

Table 3.43: Methods of `multi_ptr` class

Methods	Description
<pre>template <typename ElementType, access::address_space Space = access:: address_space::private_space> operator private_ptr<ElementType>()</pre>	Conversion operator from <code>multi_ptr<ElementType,address_space::private_space></code> to <code>private_ptr<ElementType></code> .
<pre>template <typename ElementType, access::address_space Space = access:: address_space::private_space> private_ptr<ElementType> pointer()</pre>	Pointer method that returns a <code>private_ptr<ElementType></code> from a <code>multi_ptr<ElementType,address_space::private_space></code> .
End of table	

Table 3.43: Methods of `multi_ptr` class

Non-member functions	Description
<pre>template <typename ElementType, access:: address_space Space> bool operator==(const multi_ptr<ElementType, Space>& lhs, const multi_ptr<ElementType, Space>& rhs)</pre>	Comparison operator <code>==</code> for <code>multi_ptr</code> class.
<pre>template <typename ElementType, access:: address_space Space> bool operator!=(const multi_ptr<ElementType, Space>& lhs, const multi_ptr<ElementType, Space>& rhs)</pre>	Comparison operator <code>!=</code> for <code>multi_ptr</code> class.
<pre>template <typename ElementType, access:: address_space Space> bool operator<(const multi_ptr<ElementType, Space>& lhs, const multi_ptr<ElementType, Space>& rhs)</pre>	Comparison operator <code><</code> for <code>multi_ptr</code> class.
<pre>template <typename ElementType, access:: address_space Space> bool operator>(const multi_ptr<ElementType, Space>& lhs, const multi_ptr<ElementType, Space>& rhs)</pre>	Comparison operator <code>></code> for <code>multi_ptr</code> class.
<pre>template <typename ElementType, access:: address_space Space> bool operator<=(const multi_ptr<ElementType, Space>& lhs, const multi_ptr<ElementType, Space>& rhs)</pre>	Comparison operator <code><=</code> for <code>multi_ptr</code> class.
<pre>template <typename ElementType, access:: address_space Space> bool operator>=(const multi_ptr<ElementType, Space>& lhs, const multi_ptr<ElementType, Space>& rhs)</pre>	Comparison operator <code>>=</code> for <code>multi_ptr</code> class.
<pre>template <typename ElementType, access:: address_space Space> bool operator!=(const multi_ptr<ElementType, Space>& lhs, nullptr_t rhs)</pre>	Comparison operator <code>!=</code> for <code>multi_ptr</code> class with a <code>nullptr_t</code> .
Continued on next page	

Table 3.44: Non-member functions of the `multi_ptr` class.

Non-member functions	Description
<pre>template <typename ElementType, access:: address_space Space> bool operator!=(nullptr_t lhs, const multi_ptr< ElementType, Space>& rhs)</pre>	Comparison operator != for <code>multi_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType, access:: address_space Space> bool operator==(const multi_ptr<ElementType, Space>& lhs, nullptr_t rhs)</pre>	Comparison operator == for <code>multi_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType, access:: address_space Space> bool operator==(nullptr_t lhs, const multi_ptr< ElementType, Space>& rhs)</pre>	Comparison operator == for <code>multi_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType, access:: address_space Space> bool operator>(const multi_ptr<ElementType, Space>& lhs, nullptr_t rhs)</pre>	Comparison operator > for <code>multi_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType, access:: address_space Space> bool operator>(nullptr_t lhs, const multi_ptr< ElementType, Space>& rhs)</pre>	Comparison operator > for <code>multi_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType, access:: address_space Space> bool operator<(const multi_ptr<ElementType, Space>& lhs, nullptr_t rhs)</pre>	Comparison operator < for <code>multi_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType, access:: address_space Space> bool operator<(nullptr_t lhs, const multi_ptr< ElementType, Space>& rhs)</pre>	Comparison operator < for <code>multi_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType, access:: address_space Space> bool operator>=(const multi_ptr<ElementType, Space>& lhs, nullptr_t rhs)</pre>	Comparison operator >= for <code>multi_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType, access:: address_space Space> bool operator>=(nullptr_t lhs, const multi_ptr< ElementType, Space>& rhs)</pre>	Comparison operator >= for <code>multi_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType, access:: address_space Space> bool operator<=(const multi_ptr<ElementType, Space>& lhs, nullptr_t rhs)</pre>	Comparison operator <= for <code>multi_ptr</code> class with a <code>nullptr_t</code> .
<pre>template <typename ElementType, access:: address_space Space> bool operator<=(nullptr_t lhs, const multi_ptr< ElementType, Space>& rhs)</pre>	Comparison operator <= for <code>multi_ptr</code> class with a <code>nullptr_t</code> .
End of table	

Table 3.44: Non-member functions of the `multi_ptr` class.

3.4.8 Samplers

Samplers use the `cl::sycl::sampler` type which is equivalent to the OpenCL C `cl_sampler` and `sampler_t` types. Constructors for the sampler class are listed in Table 3.45 and methods in Table 3.46.

```

1 namespace cl {
2 namespace sycl {
3 class sampler {
4 public:
5     enum class sampler_addressing_mode {
6         SYCL_SAMPLER_ADDRESS_MIRRORED_REPEAT,
7         SYCL_SAMPLER_ADDRESS_REPEAT,
8         SYCL_SAMPLER_ADDRESS_CLAMP_TO_EDGE,
9         SYCL_SAMPLER_ADDRESS_CLAMP,
10        SYCL_SAMPLER_ADDRESS_NONE
11    };
12    enum class sampler_filter_mode {
13        SYCL_SAMPLER_FILTER_NEAREST,
14        SYCL_SAMPLER_FILTER_LINEAR
15    };
16
17    sampler(bool normalized_coords, sampler_addressing_mode addressing_mode,
18           sampler_filter_mode filter_mode);
19
20    sampler(cl_sampler);
21
22    ~sampler() {}
23
24    sampler_addressing_mode get_address() const;
25
26    sampler_filter_mode get_filter() const;
27
28    cl_sampler get_opengl_sampler() const;
29 };
30 } // namespace sycl
31 } // namespace cl

```

Constructors	Description
<pre>sampler(bool normalized_coords, sampler_addressing_mode addressing_mode, sampler_filter_mode filter_mode)</pre>	<p><code>normalized_coords</code> selects whether normalized or unnormalized coordinates are used for accessing image data.</p> <p><code>addressing_mode</code> specifies how out-of-range image coordinates are handled.</p> <p><code>filter_mode</code> specifies the type of filter that must be applied when reading an image.</p>

Continued on next page

Table 3.45: Constructors for the sampler class.

Constructors	Description
<code>sampler(cl_sampler)</code>	Construct a sampler from an OpenCL sampler object. Retains a reference to the <code>cl_sampler</code> object. Caller should release the passed <code>cl_sampler</code> object when it is no longer needed.
End of table	

Table 3.45: Constructors for the `sampler` class.

Methods	Description
<code>sampler_addressing_mode get_address()const</code>	Return the addressing mode used to construct the sampler.
<code>sampler_filter_mode get_filter()const</code>	Return the filter mode used to construct the sampler.
<code>cl_sampler get_sampler()const</code>	Returns the underlying <code>cl_sampler</code> object. Retains a reference to the returned <code>cl_sampler</code> object. Caller should release it when finished.
End of table	

Table 3.46: Methods for the `sampler` class.

3.5 Expressing parallelism through kernels

3.5.1 Ranges and index space identifiers

The data parallelism of the OpenCL execution model and its exposure through SYCL requires instantiation of a parallel execution over a range of iteration space coordinates. To achieve this we expose types to define the range of execution and to identify a given execution instance's point in the iteration space.

To achieve this we expose five types: `range`, `nd_range`, `id`, `item`, `nd_item` and `group`.

When constructing ids or ranges from integers, the elements are written in row-major format.

3.5.1.1 `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. Constructors for the `range` class are described in Table 3.47, methods in Table 3.48 and global operators on ranges in Table 3.49.

An overview of the interface provided for the `range` class is the following, for the full interface please refer to B.3.

```
1 namespace cl {
```



```

2 namespace sycl {
3   template <size_t dimensions>
4   struct range {
5     range(const range<dimensions> &);
6
7     range(size_t x); // When dimensions==1
8     range(size_t x, size_t y); // When dimensions==2
9     range(size_t x, size_t y, size_t z); // When dimensions==3
10
11    size_t get(int dimension) const;
12    size_t &operator[](int dimension);
13
14    range &operator=(const range &rhs);
15    range &operator+=(const range &rhs);
16    range &operator*=(const range &rhs);
17    range &operator/=(const range &rhs);
18    range &operator%=(const range &rhs);
19    range &operator>>=(const range &rhs);
20    range &operator<<=(const range &rhs);
21    range &operator&=(const range &rhs);
22    range &operator^=(const range &rhs);
23    range &operator|=(const range &rhs);
24
25    size_t size() const;
26 };
27 } // sycl
28 } // cl

```

Constructors	Description
<code>range(const range<dimensions>&)</code>	Construct a range by deep copy from another range.
<code>range(size_t x)</code>	Construct a 1D range with value x. Only valid when for one dimension.
<code>range(size_t x, size_t y)</code>	Construct a 2D range with value x, y. Only valid when dimensions is 2.
<code>range(size_t x, size_t y, size_t z)</code>	Construct a 3D range with value x, y, z. Only valid when dimensions is 3.
End of table	

Table 3.47: Constructors for the `range` class.

Methods	Description
<code>size_t get(int dimension) const</code>	Return the value of the specified dimension of the <code>range</code> .
<code>size_t &operator[](int dimension)</code>	Return the l-value of the specified dimension of the <code>range</code> .
<code>range &operator=(const range &rhs)</code>	Assign each element of range from its corresponding element of rhs.
<code>range &operator+=(const range &rhs)</code>	Elementwise addition of rhs to the current range. Returns reference to updated range.
Continued on next page	

Table 3.48: Methods for the `range` class.

Methods	Description
<code>range &operator*=(const range &rhs)</code>	Elementwise multiplication of the current range by rhs. Returns reference to updated range.
<code>range &operator/=(const range &rhs)</code>	Elementwise division of the current range by rhs. Returns reference to updated range.
<code>range &operator%=(const range &rhs)</code>	Elementwise division of the current range by rhs, updating with the division remainder. Returns reference to updated range.
<code>range &operator>>=(const range &rhs)</code>	Elementwise arithmetic right shift of the current range by rhs. Returns reference to updated range.
<code>range &operator<<=(const range &rhs)</code>	Elementwise logical left shift of the current range by rhs. Returns reference to updated range.
<code>range &operator&=(const range &rhs)</code>	Elementwise bitwise AND of the current range by rhs. Returns reference to updated range.
<code>range &operator^=(const range &rhs)</code>	Elementwise bitwise exclusive OR of the current range by rhs. Returns reference to updated range.
<code>range &operator =(const range &rhs)</code>	Elementwise bitwise OR of the current range by rhs. Returns reference to updated range.
<code>size_t size()const;</code>	Return the size of the range computed as $\text{dimension0} * \dots * \text{dimensionN}$.
End of table	

Table 3.48: Methods for the `range` class.

Non-member functions	Description
<pre>template <size_t dimensions> bool operator ==(const range<dimensions> &a, const range<dimensions> &b)</pre>	Compare two ranges for elementwise equality. If all dimensions are equal, the ranges are equal.
<pre>template <size_t dimensions> bool operator !=(const range<dimensions> &a, const range<dimensions> &b)</pre>	Compare two ranges for elementwise inequality. If any dimension is not equal, the ranges are not equal.
<pre>template <size_t dimensions> bool operator >(const range<dimensions> &a, const range<dimensions> &b)</pre>	Compare two ranges such that a is lexically greater than b. Returns true if for any dimension n , $a[n] > b[n]$ and for all $m > n$, $a[m] == b[m]$.
<pre>template <size_t dimensions> bool operator <(const range<dimensions> &a, const range<dimensions> &b)</pre>	Compare two ranges such that a is lexically less than b. Returns true if for any dimension n , $a[n] < b[n]$ and for all $m > n$, $a[m] == b[m]$.
Continued on next page	

Table 3.49: Non-member functions for the `range` class.

Non-member functions	Description
<pre>template <size_t dimensions> bool operator >=(const range<dimensions> &a, const range<dimensions> &b)</pre>	Compare two ranges such that a is lexicographically less than or equal to b. Returns true if $a == b$ or $a > b$.
<pre>template <size_t dimensions> bool operator <=(const range<dimensions> &a, const range<dimensions> &b)</pre>	Compare two ranges such that a is lexicographically less than or equal to b. Returns true if $a == b$ or $a < b$.
<pre>template <int dimensions> range<dimensions> operator *(range<dimensions> a, range<dimensions> b)</pre>	Multiply each element of a by its respective element of b and return a range constructed from the resulting values.
<pre>template <int dimensions> range<dimensions> operator /(range<dimensions> dividend, range<dimensions> divisor)</pre>	Divide each element of dividend by its respective element in divisor and return a range constructed of the resulting value.
<pre>template <int dimensions> range<dimensions> operator +(range<dimensions> a, range<dimensions> b)</pre>	Add each element of a to its respective element of b and return a range constructed from the resulting values.
<pre>template <int dimensions> range<dimensions> operator -(range<dimensions> a, range<dimensions> b)</pre>	Subtract each element of b from its respective element of a and return a range constructed from the resulting values.
<pre>template <size_t dimensions> range<dimensions> operator %(const range<dimensions> &a, const range<dimensions> &b)</pre>	Divide each element of b from its respective element of a and return a range constructed from the remainders.
<pre>template <size_t dimensions> range<dimensions> operator <<(const range<dimensions> &a, const range<dimensions> &b)</pre>	Logically shift each element of a left by its matching element of b and return a range constructed from the shifted values.
<pre>template <size_t dimensions> range<dimensions> operator >>(const range<dimensions> &a, const range<dimensions> &b)</pre>	Arithmetically shift each element of a right by its matching element of b and return a range constructed from the shifted values.
<pre>template <size_t dimensions> range<dimensions> operator &(amp; const range<dimensions> &a, const range<dimensions> &b)</pre>	Construct a range from the bitwise AND of each element of a with the equivalent element of b.
<pre>template <size_t dimensions> range<dimensions> operator (amp; const range<dimensions> &a, const range<dimensions> &b)</pre>	Construct a range from the bitwise OR of each element of a with the equivalent element of b.

Continued on next page

Table 3.49: Non-member functions for the `range` class.

Non-member functions	Description
<pre>template <size_t dimensions> range<dimensions> operator &&(const range<dimensions> &a, const range<dimensions> &b)</pre>	Construct a range from the logical AND of each element of a with the equivalent element of b.
<pre>template <size_t dimensions> range<dimensions> operator (const range<dimensions> &a, const range<dimensions> &b)</pre>	Construct a range from the logical OR of each element of a with the equivalent element of b.
<pre>template <size_t dimensions> range<dimensions> operator ^(const range<dimensions> &a, const range<dimensions> &b)</pre>	Construct a range from the bitwise XOR of each element of a with the equivalent element of b.
<pre>template <size_t dimensions> range<dimensions> operator *(const size_t &a, const range<dimensions> &b)</pre>	Construct a range from the multiplication of each element of b with a.
<pre>template <size_t dimensions> range<dimensions> operator *(const range<dimensions> &a, const size_t &b)</pre>	Construct a range from the multiplication of each element of a with b.
<pre>template <size_t dimensions> range<dimensions> operator /(const size_t &a, const range<dimensions> &b)</pre>	Construct a range from the division of each element of a by b.
<pre>template <size_t dimensions> range<dimensions> operator /(const range<dimensions> &a, const size_t &b)</pre>	Construct a range from the division of each element of a with b.
<pre>template <size_t dimensions> range<dimensions> operator +(const size_t &a, const range<dimensions> &b)</pre>	Construct a range from the addition of each element of b with a.
<pre>template <size_t dimensions> range<dimensions> operator +(const range<dimensions> &a, const size_t &b)</pre>	Construct a range from the addition of each element of a with b.
<pre>template <size_t dimensions> range<dimensions> operator -(const size_t &a, const range<dimensions> &b)</pre>	Construct a range from the subtraction of each element of b with a.
<pre>template <size_t dimensions> range<dimensions> operator -(const range<dimensions> &a, const size_t &b)</pre>	Construct a range from the subtraction of each element of a with b.
Continued on next page	

Table 3.49: Non-member functions for the `range` class.

Non-member functions	Description
<pre>template <size_t dimensions> range<dimensions> operator %(const size_t &a, const range<dimensions> &b)</pre>	Construct a range from the modulo of each element of b with a.
<pre>template <size_t dimensions> range<dimensions> operator %(const range<dimensions> &a, const size_t &b)</pre>	Construct a range from the modulo of each element of a with b.
<pre>template <size_t dimensions> range<dimensions> operator <<(const size_t &a, const range<dimensions> &b)</pre>	Construct a range from shifting left of each element of b with a.
<pre>template <size_t dimensions> range<dimensions> operator <<(const range<dimensions> &a, const size_t &b)</pre>	Construct a range from shifting left of each element of a with b.
<pre>template <size_t dimensions> range<dimensions> operator >>(const size_t &a, const range<dimensions> &b)</pre>	Construct a range from shifting right of each element of b with a.
<pre>template <size_t dimensions> range<dimensions> operator >>(const range<dimensions> &a, const size_t &b)</pre>	Construct a range from shifting right of each element of a with b.
End of table	

Table 3.49: Non-member functions for the `range` class.

3.5.1.2 `nd_range` class

```

1 namespace cl {
2 namespace sycl {
3   template <int dimensions>
4   struct nd_range {
5     nd_range(const nd_range<dimensions> &);
6
7     nd_range(range<dims> globalSize, range<dims> localSize,
8             id<dims> offset = id<dims>());
9
10    range<dims> get_global() const;
11    range<dims> get_local() const;
12    range<dims> get_group() const;
13    id<dims> get_offset() const;
14  };
15 } // namespace sycl
16 } // namespace cl
```

`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. Constructors for the `nd_range` class are described in Table 3.50 and methods in Table 3.51.

Constructors	Description
<code>nd_range(const nd_range<dimensions> &)</code>	Construct an <code>nd_range</code> by deep copy from another <code>ndrange</code> .
<code>nd_range<dimensions>(range<dimensions> globalSize, range<dimensions> localSize) id<dimensions> offset = id<dimensions>())</code>	Construct an <code>nd_range</code> from the local and global constituent ranges as well as an optional offset. If the offset is not provided it will default to no offset.
End of table	

Table 3.50: Constructors for the `nd_range` class.

Methods	Description
<code>range<dimensions> get_global()const</code>	Return the constituent global range.
<code>range<dimensions> get_local()const</code>	Return the constituent local range.
<code>range<dimensions> get_group()const</code>	Return a range representing the number of groups in each dimension. This range would result from <code>globalSize/localSize</code> as provided on construction.
<code>id<dimensions> get_offset()const</code>	Return the constituent offset.
End of table	

Table 3.51: Methods for the `nd_range` class.

3.5.1.3 `id` class

`id<int dimensions>` is a vector of dimensions that is used to represent an *index* 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`. Constructors for the `id` class are described in Table 3.52, methods in Table 3.53 and global operators on `ids` in Table 3.54

An overview of the interface provided for the `id` class is the following, for the full interface please refer to B.4.

```

1 namespace cl {
2 namespace sycl {
3 template <size_t dimensions>
4 struct id {
5     id(size_t x); // When dimensions==1
6     id(size_t x, size_t y); // When dimensions==2
7     id(size_t x, size_t y, size_t z); // When dimensions==3
8     id(const id<dimensions> & rhs);
9     id(const range<dimensions> & rangeSize);
10    id(const item<dimensions> & rhs);
11
12    size_t get(int dimension) const;
13    size_t &operator[](int dimension);
14    operator size_t(); // When dimensions==1

```

```

15
16   id &operator=(const id & rhs);
17   id &operator+=(const id & rhs);
18   id &operator*=(const id & rhs);
19   id &operator/=(const id & rhs);
20   id &operator%=(const id & rhs);
21   id &operator>>=(const id & rhs);
22   id &operator<<=(const id & rhs);
23   id &operator&=(const id & rhs);
24   id &operator^=(const id & rhs);
25   id &operator|=(const id & rhs);
26 };
27 } // namespace sycl
28 } // namespace cl

```

Constructors	Description
<code>id(size_t x)</code>	Construct a 1D id with value x. Only valid when dimensions is 1.
<code>id(size_t x, size_t y)</code>	Construct a 1D id with value x, y. Only valid when dimensions is 2.
<code>id(size_t x, size_t y, size_t z)</code>	Construct a 1D id with value x, y, z. Only valid when dimensions is 3.
<code>id(const id &)</code>	Construct an id by deep copy.
<code>id(const range &r)</code>	Construct an id from the dimensions of a range.
<code>id(const item &it)</code>	Construct an id from <code>it.get_global_id()</code> .
End of table	

Table 3.52: Constructors for the `id` class.

Methods	Description
<code>size_t get(int dimension) const</code>	Return the value of the <code>id</code> for dimension <code>dimension</code> .
<code>size_t &operator[](int dimension) const</code>	Return a reference to the requested dimension of the id object.
<code>operator size_t()</code>	Conversion operator so that a <code>id<1></code> can be used as a plain <code>size_t</code> .
<code>id &operator=(const id &rhs)</code>	Elementwise assignment of id rhs to current id.
<code>id &operator+=(const id &rhs)</code>	Elementwise addition of rhs to the current id. Returns reference to updated id.
<code>id &operator*=(const id &rhs)</code>	Elementwise multiplication of the current id by rhs. Returns reference to updated id.
<code>id &operator/=(const id &rhs)</code>	Elementwise division of the current id by rhs. Returns reference to updated id.
<code>id &operator%=(const id &rhs)</code>	Elementwise division of the current id by rhs, updating with the division remainder. Returns reference to updated id.
Continued on next page	

Table 3.53: Methods for the `id` class.

Methods	Description
<code>id &operator>>=(const id &rhs)</code>	Elementwise arithmetic right shift of the current id by rhs. Returns reference to updated id.
<code>id &operator<<=(const id &rhs)</code>	Elementwise logical left shift of the current id by rhs. Returns reference to updated id.
<code>id &operator&=(const id &rhs)</code>	Elementwise bitwise AND of the current id by rhs. Returns reference to updated id.
<code>id &operator^=(const id &rhs)</code>	Elementwise bitwise exclusive OR of the current id by rhs. Returns reference to updated id.
<code>id &operator =(const id &rhs)</code>	Elementwise bitwise OR of the current id by rhs. Returns reference to updated id.
End of table	

Table 3.53: Methods for the `id` class.

Non-member functions	Description
<pre>template <size_t dimensions> bool operator ==(const id<dimensions> &a, const id<dimensions> &b)</pre>	Compare two ids for elementwise equality. If all dimensions are equal, the ids are equal.
<pre>template <size_t dimensions> bool operator !=(const id<dimensions> &a, const id<dimensions> &b)</pre>	Compare two ids for elementwise inequality. If any dimension is not equal, the ids are not equal.
<pre>template <size_t dimensions> bool operator >(const id<dimensions> &a, const id<dimensions> &b)</pre>	Compare two ids such that a is lexically greater than b. Returns true if for any dimension n, $a[n] > b[n]$ and for all $m > n$, $a[m] == b[m]$.
<pre>template <size_t dimensions> bool operator <(const id<dimensions> &a, const id<dimensions> &b)</pre>	Compare two ids such that a is lexically less than b. Returns true if for any dimension n, $a[n] < b[n]$ and for all $m > n$, $a[m] == b[m]$.
<pre>template <size_t dimensions> bool operator >=(const id<dimensions> &a, const id<dimensions> &b)</pre>	Compare two ids such that a is lexically less than or equal to b. Returns true if $a == b$ or $a > b$.
<pre>template <size_t dimensions> bool operator <=(const id<dimensions> &a, const id<dimensions> &b)</pre>	Compare two ids such that a is lexically less than or equal to b. Returns true if $a == b$ or $a < b$.
<pre>template <size_t dimensions> id<dimensions> operator *(const id<dimensions> &a, const id<dimensions> &b)</pre>	Multiply each element of a by its respective element of b and return an id constructed from the resulting values.
Continued on next page	

Table 3.54: Non-member functions for the `id` class.

Non-member functions	Description
<pre>template <size_t dimensions> id<dimensions> operator /(const id<dimensions> &dividend, const id<dimensions> &divisor)</pre>	Divide each element of dividend by its respective element in divisor and return an id constructed of the resulting value.
<pre>template <size_t dimensions> id<dimensions> operator +(const id<dimensions> &a, const id<dimensions> &b)</pre>	Add each element of a to its respective element of b and return an id constructed from the resulting values.
<pre>template <size_t dimensions> id<dimensions> operator -(const id<dimensions> &a, const id<dimensions> &b)</pre>	Subtract each element of b from its respective element of a and return an id constructed from the resulting values.
<pre>template <size_t dimensions> id<dimensions> operator %(const id<dimensions> &a, const id<dimensions> &b)</pre>	Divide each element of b from its respective element of a and return an id constructed from the remainders.
<pre>template <size_t dimensions> id<dimensions> operator <<(const id<dimensions> &a, const id<dimensions> &b)</pre>	Logically shift each element of a left by its matching element of b and return an id constructed from the shifted values.
<pre>template <size_t dimensions> id<dimensions> operator >>(const id<dimensions> &a, const id<dimensions> &b)</pre>	Arithmetically shift each element of a right by its matching element of b and return an id constructed from the shifted values.
<pre>template <size_t dimensions> id<dimensions> operator &(const id<dimensions> &a, const id<dimensions> &b)</pre>	Construct an id from the bitwise AND of each element of a with the equivalent element of b.
<pre>template <size_t dimensions> id<dimensions> operator (const id<dimensions> &a, const id<dimensions> &b)</pre>	Construct an id from the bitwise OR of each element of a with the equivalent element of b.
<pre>template <size_t dimensions> id<dimensions> operator ^(const id<dimensions> &a, const id<dimensions> &b)</pre>	Construct an id from the bitwise XOR of each element of a with the equivalent element of b.
<pre>template <size_t dimensions> id<dimensions> operator &&(const id<dimensions> &a, const id<dimensions> &b)</pre>	Construct an id from the logical AND of each element of a with the equivalent element of b.
<pre>template <size_t dimensions> id<dimensions> operator (const id<dimensions> &a, const id<dimensions> &b)</pre>	Construct an id from the logical OR of each element of a with the equivalent element of b.
Continued on next page	

Table 3.54: Non-member functions for the `id` class.

Non-member functions	Description
<pre>template <size_t dimensions> id<dimensions> operator *(const size_t &a, const id<dimensions> &b)</pre>	Construct an id from the multiplication of a with each element of a.
<pre>template <size_t dimensions> id<dimensions> operator *(const id<dimensions> &a, const size_t &b)</pre>	Construct an id from the multiplication of each element of a with b.
<pre>template <size_t dimensions> id<dimensions> operator +(const size_t &a, const id<dimensions> &b)</pre>	Construct an id from the addition of a with each equivalent element of b.
<pre>template <size_t dimensions> id<dimensions> operator +(const id<dimensions> &a, const size_t &b)</pre>	Construct an id from the addition of each element of a with b.
<pre>template <size_t dimensions> id<dimensions> operator -(const size_t &a, const id<dimensions> &b)</pre>	Construct an id from the subtraction of a with each equivalent element of b.
<pre>template <size_t dimensions> id<dimensions> operator -(const id<dimensions> &a, const size_t &b)</pre>	Construct an id from the subtraction of each element of a with b.
<pre>template <size_t dimensions> id<dimensions> operator /(const size_t &a, const id<dimensions> &b)</pre>	Construct an id from the division of a with each equivalent element of b.
<pre>template <size_t dimensions> id<dimensions> operator /(const id<dimensions> &a, const size_t &b)</pre>	Construct an id from the division of each element of a with b.
<pre>template <size_t dimensions> id<dimensions> operator %(const size_t &a, const id<dimensions> &b)</pre>	Construct an id from the modulo of a with each equivalent element of b.
<pre>template <size_t dimensions> id<dimensions> operator %(const id<dimensions> &a, const size_t &b)</pre>	Construct an id from the modulo of each element of a with b.
<pre>template <size_t dimensions> id<dimensions> operator <<(const size_t &a, const id<dimensions> &b)</pre>	Construct an id from shifting left a with each equivalent element of b.

Continued on next page

Table 3.54: Non-member functions for the `id` class.

Non-member functions	Description
<pre>template <size_t dimensions> id<dimensions> operator <<(const id<dimensions> &a, const size_t &b)</pre>	Construct an id from shifting left each element of a with b.
<pre>template <size_t dimensions> id<dimensions> operator >>(const size_t &a, const id<dimensions> &b)</pre>	Construct an id from shifting right a with each equivalent element of b.
<pre>template <size_t dimensions> id<dimensions> operator >>(const id<dimensions> &a, const size_t &b)</pre>	Construct an id from shifting right each element of a with b.
End of table	

Table 3.54: Non-member functions for the `id` class.

3.5.1.4 `item` class

`item<int dimensions>` identifies an instance of the functor executing at each point in a `range<>` passed to a `parallel_for` call, or passed to a `parallel_for_work_item` call. It encapsulates enough information to identify the work-item's global or local ID, the range of possible values, and 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 functor. Methods for the `item<>` class are described in Table 3.55.

```

1 namespace cl {
2 namespace sycl {
3 template <int dimensions>
4 struct item {
5     item() = delete;
6
7     item(const item &rhs) = default;
8
9     id<dimensions> get() const;
10
11     size_t get(int dimension) const;
12
13     size_t &operator[](int dimension);
14
15     range<dimensions> get_range() const;
16
17     id<dimensions> get_offset() const;
18
19     size_t get_linear_id() const;
20
21     operator id<dimensions>();
22 };
23 } // namespace sycl
24 } // namespace cl
```

Methods	Description
<code>id<dimensions> get()const</code>	Return the constituent local or global <code>id<></code> representing the work-item's position in the iteration space.
<code>size_t get(int dimension)const</code>	Return the requested dimension of the constituent <code>id<></code> representing the work-item's position in the iteration space.
<code>size_t &operator[](int dimension)</code>	Return the constituent <code>id<></code> l-value representing the work-item's position in the iteration space in the given dimension.
<code>range<dimensions> get_range()const</code>	Returns a <code>range<></code> representing the dimensions of the range of possible values of the <code>item</code> .
<code>id<dimensions> get_offset()const</code>	Returns an <code>id<></code> representing the n -dimensional offset provided to the <code>parallel_for</code> 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 representing a local range of where no offset was passed this will always return an <code>id</code> of all 0 values.
<code>size_t get_linear_id()const</code>	Return the linearized ID in the <code>item</code> 's range. Computed as the flattened ID after the offset is subtracted.
End of table	

Table 3.55: Methods for the `item` class.

3.5.1.5 `nd_item` class

`nd_item<int dimensions>` identifies an instance of the functor executing at each point in an `nd_range<int dimensions>` passed to a `parallel_for_ndrange` call. It encapsulates enough information to identify the work-item's local and global IDs, the work-groups ID and also provides barrier functionality to synchronize work-items. Instances of the `nd_item<int dimensions>` class are not user-constructible and are passed by the runtime to each instance of the functor. Methods for the `nd_item<int dimensions>` class are described in Table 3.56.

```

1 namespace cl {
2 namespace sycl {
3 template <int dimensions>
4 struct nd_item {
5     nd_item() = delete;
6
7     nd_item(const nd_item &) = default;
8
9     id<dimensions> get_global() const;
10
11     size_t get_global(int) const;
12
13     size_t get_global_linear_id() const;
14

```

```

15  id<dimensions> get_local() const;
16
17  size_t get_local(int) const;
18
19  size_t get_local_linear_id() const;
20
21  group<dimensions> get_group() const;
22
23  size_t get_group(int) const;
24
25  size_t get_group_linear_id() const;
26
27  int get_num_groups(int) const;
28
29  id<dimensions> get_num_groups() const;
30
31  range<dimensions> get_global_range() const;
32
33  range<dimensions> get_local_range() const;
34
35  id<dimensions> get_offset() const;
36
37  nd_range<dimensions> get_nd_range() const;
38
39  void barrier(access::fence_space flag = access::global_and_local) const;
40 };
41 } // namespace sycl
42 } // namespace cl

```

Methods	Description
<code>id<dimensions> get_global()const</code>	Return the constituent global <code>id</code> representing the work-item's position in the global iteration space.
<code>size_t get_global(int dimension)const</code>	Return the constituent element of the global <code>id</code> representing the work-item's position in the global iteration space in the given dimension.
<code>id<dimensions> get_global_linear_id()const</code>	Return the flattened id of the current work-item after subtracting the offset.
<code>id<dimensions> get_local()const</code>	Return the constituent local <code>id</code> representing the work-item's position within the current work-group.
<code>size_t get_local(int dimension)const</code>	Return the constituent element of the local <code>id</code> representing the work-item's position within the current work-group in the given dimension.
<code>id<dimensions> get_local_linear_id()const</code>	Return the flattened id of the current work-item within the current work-group.
<code>group<dimensions> get_group()const</code>	Return the constituent group <code>group</code> representing the work-group's position within the overall <code>nd_range</code> .

Continued on next page

Table 3.56: Methods for the `nd_item` class.

Methods	Description
<code>size_t get_group(int dimension) const</code>	Return the constituent element of the group <code>id</code> representing the work-group's position within the overall <code>nd_range</code> in the given dimension.
<code>size_t get_group_linear_id() const</code>	Return the flattened id of the current work-group.
<code>int get_num_groups(int dimension) const</code>	Return the number of groups for dimension in the <code>nd_range</code> .
<code>id<dimensions> get_num_groups() const</code>	Returns the number of groups in the <code>nd_range</code> .
<code>range<dimensions> get_global_range() const</code>	Returns a <code>range<></code> representing the dimensions of the <code>nd_range<></code>
<code>range<dimensions> get_local_range() const</code>	Returns a <code>range<></code> representing the dimensions of the current work-group.
<code>id<dimensions> get_offset() const</code>	Returns an <code>id<></code> representing the n-dimensional offset provided to the constructor of the <code>nd_range<></code> and that is added by the runtime to the global-ID of each work-item.
<code>nd_range<dimensions> get_nd_range() const</code>	Returns the <code>nd_range<></code> of the current execution.
<code>void barrier(access::fence_space flag= access::global_and_local) const</code>	Executes a barrier with memory ordering on the local address space, global address space or both based on the value of <code>flag</code> . The current work-item will wait at the barrier until all work-items in the current work-group have reached the barrier. In addition the barrier performs a fence operation ensuring that all memory accesses in the specified address space issued before the barrier complete before those issued after the barrier.
End of table	

Table 3.56: Methods for the `nd_item` class.

3.5.1.6 group class

The `group<int dimensions>` is passed to each instance of the functor execution a `parallel_for_workgroup` in a hierarchical parallel execution. The group encapsulates all functionality required to represent a particular group within a parallel execution. It is not user-constructable. Methods for the `group<>` class are described in Table 3.57.

The local range stored in the group class will be 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 dimensions of the `parallel_for_work_item` loops.

```
1 namespace cl {
```

```

2 namespace sycl {
3 template <int dimensions>
4 struct group {
5     group(const group &rhs) = default;
6
7     id<dimensions> get() const;
8
9     size_t get(int dimension) const;
10
11    range<dimensions> get_global_range() const;
12
13    size_t get_global_range(int) const;
14
15    range<dimensions> get_group_range() const;
16
17    size_t get_group_range(int) const;
18
19    size_t operator[](int) const;
20
21    size_t get_linear() const;
22 };
23 } // sycl
24 } // cl

```

Methods	Description
<code>id<dimensions> get()</code>	Return an <code>id</code> representing the index of the group within the <code>nd_range</code> for every dimension.
<code>size_t get(int dimension) const</code>	Return the index of the group in the given dimension.
<code>range<dimensions> get_global_range()</code>	Return the constituent global range.
<code>size_t get_global_range(int dimension)</code>	Return element dimension from the constituent global range.
<code>range<dimensions> get_group_range()</code>	Return a <code>range<></code> representing the dimensions of the current group. This local range may have been provided by the programmer, or chosen by the runtime.
<code>size_t get_group_range(int dimension)</code>	Return element dimension from the constituent group range.
<code>size_t operator[](int dimension) const</code>	Return the index of the group in the given dimension within the <code>nd_range<></code> .
<code>size_t get_linear() const</code>	Get a linearized version of the group ID.
End of table	

Table 3.57: Methods for the `group` class.

3.5.2 Command group scope

A *command group scope* in SYCL as it is defined in Section 2.3.1 consists of a kernel and all the commands for queued data transfers in order for the kernel’s execution to be successful. In SYCL, as it is in OpenCL,

a kernel needs to be enqueued and there are a lot of commands that need to be enqueued as well for making the data available to the kernel. The commands that enqueue a kernel and all the relevant data transfers to be enqueued for it, form the *command group functor*. The command group functor takes as a parameter an instance of the command group [handler 3.5.3.4](#) class which encapsulates all the methods executed in the command group scope. 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 functor and the [handler](#) class serves as an interface for the encapsulation of *command group scope*. A kernel is defined in a command group either as a functor object or as a lambda function. All the device data accesses are defined inside this group and any transfers are managed by the system. The rules for the data transfers regarding device and host data accesses are better described in the data management section ([3.4](#)), where buffers ([3.4.2](#)) and accessor ([3.4.6](#)) classes are described.

It is possible to obtain events for the start of the command group functor, 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 due to the fact that it is not rigidly specified which memory data are stored on kernel completion. The events are provided at the submission of the command group functor at the queue to be executed on.

It is possible for a command group functor 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 undefined in the specification.

The command group [handler](#) class provides the interface for all of the methods that are able to be executed inside the command group scope and also it is provided as a scoped object to all of the data access requests. The command group [handler](#) class provides the interface in which every command of the command group scope will be submitted to a queue.

3.5.3 SYCL functions for invoking kernels

Kernels can be invoked as *single tasks*, basic *data-parallel kernels*, OpenCL-style *NDRanges* in *work-groups*, or SYCL *hierarchical parallelism*.

Each function takes a kernel name template parameter. The kernel name must be a datatype that is unique for each kernel invocation. If a kernel is a functor, and its type is globally visible, then the kernel's functor type will be automatically used as the kernel name and so the user does not need to supply a name. If the kernel function is a C++11 lambda function, then the user must manually provide a kernel name to enable linking between host and device code to occur.

All the functions for invoking kernels are methods of the command group `cl::sycl::handler` class [3.5.3.4](#), which is used to encapsulate all the methods provided in a command group scope.

3.5.3.1 `single_task` invoke

SYCL provides a simple interface to enqueue a kernel that will be sequentially executed on an OpenCL device. Only one instance of the kernel will be executed. This interface is useful a primitive for more complicated parallel algorithms, as it can easily create a chain of sequential tasks on an OpenCL device with each of them managing

its own data transfers.

This function can only be called inside a command group using the command group `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.

```
1 auto command_group = [&](handler & cgh) {
2   cgh.single_task<class kernel_name>(
3     [=] () {
4       // [kernel code]
5     });
6 };
```

For single tasks, the kernel method takes no parameters, as there is no need for indexing classes in a unary index space.

3.5.3.2 `parallel_for` invoke

The `parallel_for` interface offers the ability to the SYCL users to declare a kernel and enqueue it as a parallel execution over a range of instances. There are three variations to the `parallel_for` interface and they depend on the index space that the developer would like the kernels to be executed on and the feature set available in those kernels.

In the simplest case, the developer only needs to provide the number of work-items the kernel will use in total and the system will use the best range available to enqueue it on a device. In this case the developer, may only need to know the index over the total *range* that he has provided, by providing the number of work-items that will be executing on. This type of kernels will be using the `parallel_for` invocation with a `range` type to provide the range of the execution and an `id` to provide the index within that range. Whether it is a lambda function or a kernel functor the parameter to the invocation function need to be `id`.

An example of a `parallel_for` using a lambda function for a kernel invocation in this case of `parallel_for` is the following.

```
1 class MyKernel;
2
3 myQueue.submit( [&](handler & cmdgroup)
4 {
5   auto acc=myBuffer.get_access<read_write>();
6
7   cmdgroup.parallel_for<class MyKernel>(range<1>(workItemNo),
8     [=] (id<1> index)
9     {
10      acc[index] = 42.0f;
11    });
12 });
```

Another case, which is based on this very basic `parallel_for`, is the case where the developer would like to let the runtime choose the index space that is matching best the *range* provided but would like to use information given the scheduled interface instead of the general interface. This is enabled by using the class `item` as an indexing

class in the kernel, and of course that would mean that the kernel invocation would match the `range` with the `item` parameter to the kernel.

```
1 class MyKernel;
2
3 myQueue.submit([& (handler & cmdgroup)
4 {
5     auto acc=myBuffer.get_access<read_write>();
6
7     cmdgroup.parallel_for<class MyKernel>(range<1>(workItemNo),
8         [=] (item<1> myItem)
9         {
10             size_t index = item.get_global();
11             acc[index] = 42.0f;
12         });
13 });
```

Local accessors are disallowed for the basic `parallel_for` invocations described above.

The following two examples show how a kernel functor can be launched over a 3D grid, 3 elements in each dimension. In the first case work-item IDs range from 0 to 2 inclusive, in the second case work-item IDs they run from 1 to 3.

```
1 auto command_group = [& (handler & cgh) {
2     cgh.parallel_for<class example_kernel1>(
3         range<3>(3,3,3), // global range
4         [=] (item<3> it) {
5             //[kernel code]
6         });
7 };
8 auto command_group2 = [& (handler & cgh) {
9     cgh.parallel_for<class example_kernel2>(
10        range<3>(3,3,3), // global range
11        id<3>(1,1,1), // offset
12        [=] (item<3> it) {
13            //[kernel code]
14        });
15 };
```

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, which provides all the functionality of OpenCL for an `NDRange`. 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 enqueueing 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 functor or lambda is passed an `nd_item<i>` instance making all the information available as well as barrier primitives to synchronize the work-items in the 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 sixteen work-groups. Each group of work-items synchronizes with a barrier.

```
1 auto command_group = [& (handler& cgh) {
```

```

2   cgh.parallel_for<class example_kernel>(
3       nd_range(range(4, 4, 4), range(2, 2, 2)), [=](nd_item<3> item) {
4           //[kernel code]
5           // Internal synchronization
6           item.barrier(access::fence_space::global);
7           //[kernel code]
8       });
9 };

```

Optionally, in any of these variations of `parallel_for` invocations, the developer may also pass an offset. An offset is an instance of the `id` class added to the identifier for each point in the range.

In all of these case the underlying `nd_range` will be created and the kernel defined as a lambda or as a kernel functor will be created and enqueued as part of the command group scope.

3.5.3.3 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. `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 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 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 directly 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 very 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:

```

1   template <typename T, int Dimensions>
2   class private_memory {
3   public:
4       // Construct based directly off the number of work-items
5       private_memory(const group<Dimensions> &);
6
7       // Access the instance for the current work-item
8       T &operator()(const item<Dimensions> &id);
9   };

```

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.

Private memory will be reused modulo the underlying work-group size in each dimension. For example, a 2x2 work-group will allocate 4 private memory variables and repeat them three times horizontally and twice vertically in a 6x4 `parallel_for_work_item` loop.

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 barrier in the earlier example.

```

1 auto command_group = [&](handler & cgh) {
2   // Issue 8 work-groups of 8 work-items each
3   cgh.parallel_for_work_group<class example_kernel>(
4     range<3>(2, 2, 2), range<3>(2, 2, 2), [=](group<3> myGroup) {
5
6     //[workgroup code]
7     int myLocal; // this variable shared between workitems
8     // this variable will be instantiated for each work-item separately
9     private_memory<int> myPrivate(myGroup);
10
11    // Issue parallel sets of work-items each sized using the runtime default
12    parallel_for_work_item(myGroup, [=](item<3> myItem) {
13      //[work-item code]
14      myPrivate(myItem) = 0;
15    });
16
17    // Carry private value across loops
18    parallel_for_work_item(myGroup, [=](item<3> myItem) {
19      //[work-item code]
20      output[myGroup.get_local_range()*myGroup.get()+myItem] =
21        myPrivate(myItem);
22    });
23    //[workgroup code]
24  });
25 };

```

It is valid to use more flexible dimensions of the work-item loops. In the following example we issue 8 work-groups 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.

```

1 auto flexible_command_group = [&](handler & cgh) {
2   // Issue 1000 work-groups of 8 work-items each
3   cgh.parallel_for_work_group<class example_kernel>(
4     range<3>(2, 2, 2), [=](group<3> myGroup) {
5
6     // Launch a set of 8 work-items as requested in the parallel_for_work_group launch
7     parallel_for_work_item(myGroup, [=](item<3> myItem) {
8       //[work-item code]

```

```

9     });
10    // Launch 512 iterations that will map to the underlying 8
11    parallel_for_work_item(myGroup, range<3>(8, 8, 8), [=](item<3> myItem) {
12        //[work-item code]
13    });
14    //[workgroup code]
15 });
16 });

```

This interface offers a more intuitive way to 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_workgroup` 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 a host execution may be more efficient.

3.5.3.4 Command group handler class

A command group `handler` object can only be constructed by the SYCL runtime. All of the accessors defined in the command group scope take as a parameter an instance of the command group `handler` and all the kernel invocation functions are methods of this class.

The constructors and methods of the command group `handler` class are listed in Tables 3.58 and 3.59.

```

1 namespace cl {
2 namespace sycl {
3
4 class handler_event {
5 public:
6     event get_kernel() const;
7     event get_complete() const;
8     event get_end() const;
9 }
10
11 class handler {
12 private:
13     // implementation defined constructor
14     handler(__unspecified__);
15
16 public:
17     handler(const handler& rhs);
18
19     void set_arg(int arg_index, accessor acc_obj);
20
21     template <typename T>
22     void set_arg(int arg_index, T scalar_value);
23
24     //In the case of a functor with a globally visible name
25     //the template parameter:"typename kernelName" can be omitted
26     //and the kernelType can be used instead.
27     template <typename KernelName, class KernelType>
28     void single_task(KernelType);
29
30     //In the case of a functor with a globally visible name

```

```

31 //the template parameter:"typename kernelName" can be omitted
32 //and the kernelType can be used instead.
33 template <typename KernelName, class KernelType>
34 void parallel_for(range<dimensions> numWorkItems, KernelType);
35
36 //In the case of a functor with a globally visible name
37 //the template parameter:"typename kernelName" can be omitted
38 //and the kernelType can be used instead.
39 template <typename KernelName, class KernelType>
40 void parallel_for(range<dimensions> numWorkItems,
41                 id<dimensions> workItemOffset, KernelType);
42
43 //In the case of a functor with a globally visible name
44 //the template parameter:"typename kernelName" can be omitted
45 //and the kernelType can be used instead.
46 template <typename KernelName, class KernelType>
47 void parallel_for(nd_range<dimensions> executionRange, KernelType);
48
49 //In the case of a functor with a globally visible name
50 //the template parameter:"typename kernelName" can be omitted
51 //and the kernelType can be used instead.
52 template <typename KernelName, class KernelType>
53 void parallel_for(nd_range<dimensions> numWorkItems,
54                 id<dimensions> workItemOffset, KernelType);
55
56 //In the case of a functor with a globally visible name
57 //the template parameter:"typename kernelName" can be omitted
58 //and the kernelType can be used instead.
59 template <class KernelName, class WorkgroupFunctionType>
60 void parallel_for_work_group(range<dimensions> numWorkGroups,
61                             WorkgroupFunctionType);
62
63 //In the case of a functor with a globally visible name
64 //the template parameter:"typename kernelName" can be omitted
65 //and the kernelType can be used instead.
66 template <class KernelName, class WorkgroupFunctionType>
67 void parallel_for_work_group(range<dimensions> numWorkGroups,
68                             range<dimensions> workGroupSize,
69                             WorkgroupFunctionType);
70
71 void single_task(kernel syclKernel);
72
73 void parallel_for(range<dimensions> numWorkItems, kernel syclKernel);
74
75 void parallel_for(nd_range<dimensions> ndRange, kernel syclKernel);
76
77 };
78 } // namespace sycl
79 } // namespace cl

```

Constructors	Description
<code>handler(const handler &rhs)</code>	Copy constructor of a command group <i>handler</i> .

Table 3.58: Constructors for the command group `handler` class

Methods	Description
<code>void set_arg(int index, accessor & accObj)</code>	Set kernel args for an OpenCL kernel which is used through the SYCL/OpenCL interop interface. The index value specifies which parameter of the OpenCL kernel is being set and the accessor object, which OpenCL buffer or image is going to be given as kernel argument.
<code>template <typename T> void set_arg(int index, accessor & accObj)</code>	Set kernel args for an OpenCL kernel which is used through the SYCL/OpenCL interoperability interface. The index value specifies which parameter of the OpenCL kernel is being set and the accessor object, which OpenCL buffer or image is going to be given as kernel argument.
<code>template <typename KernelName, class KernelType> void single_task(KernelType)</code>	Kernel invocation method of a kernel defined as a lambda or functor. If it is a lambda function or the functor type is globally visible there is no need for the developer to provide a kernel name type (<code>typename KernelName</code>) for it, as described in 3.5.3
<code>template <typename KernelName, class KernelType> void parallel_for(range<dimensions> numWorkItems, KernelType)</code>	Kernel invocation method of a kernel defined as a lambda or functor, for the specified range and given an id or item for indexing in the indexing space defined by range. If it is a lambda function or the if the functor type is globally visible there is no need for the developer to provide a kernel name type (<code>typename KernelName</code>) for it, as described in detail in 3.5.3
<code>template <typename KernelName, class KernelType> void parallel_for(range<dimensions> numWorkItems, id<dimensions> workItemOffset, KernelType)</code>	Kernel invocation method of a kernel defined as a lambda or functor, for the specified range and offset and given an id or item for indexing in the indexing space defined by range. If it is a lambda function or the if the functor type is globally visible there is no need for the developer to provide a kernel name type (<code>typename KernelName</code>) for it, as described in detail in 3.5.3
<code>template <typename KernelName, class KernelType> void parallel_for(nd_range<dimensions> executionRange, KernelType)</code>	Kernel invocation method of a kernel defined as a lambda or functor, for the specified <code>nd_range</code> and given an <code>nd_item</code> for indexing in the indexing space defined by the <code>nd_range</code> . If it is a lambda function or the if the functor type is globally visible there is no need for the developer to provide a kernel name type (<code>typename KernelName</code>) for it, as described in detail in 3.5.3

Continued on next page

Table 3.59: Methods for the command group `handler` class

Methods	Description
<pre>template <typename KernelName, class KernelType> void parallel_for(nd_range<dimensions> numWorkItems id<dimensions> workItemOffset, KernelType)</pre>	Kernel invocation method of a kernel defined as a lambda or functor, for the specified <code>nd_range</code> and given an <code>nd_item</code> for indexing in the indexing space defined by the <code>nd_range</code> . If it is a lambda function or the if the functor type is globally visible there is no need for the developer to provide a kernel name type (<code>typename KernelName</code>) for it, as described in detail in 3.5.3
<pre>template <class KernelName, class WorkgroupFunctionType> void parallel_for_work_group(range<dimensions> numWorkGroups, WorkgroupFunctionType)</pre>	Hierarchical kernel invocation method of a kernel defined as a lambda encoding the body of each work-group to launch. May contain multiple kernel built-in <code>parallel_for_work_item</code> functions representing the execution on each work-item. Launches <code>num_work_groups</code> work-groups of runtime-defined size. Described in detail in 3.5.3 .
<pre>template <class KernelName, class WorkgroupFunctionType> void parallel_for_work_group(range<dimensions> numWorkGroups, range<dimensions> workGroupSize, WorkgroupFunctionType)</pre>	Hierarchical kernel invocation method of a kernel defined as a lambda encoding the body of each work-group to launch. May contain multiple kernel built-in <code>parallel_for_work_item</code> functions representing the execution on each work-item. Launches <code>num_work_groups</code> work-groups of <code>work_group_size</code> work-items each. Described in detail in 3.5.3 .
<pre>void single_task(kernel syclKernel)</pre>	Kernel invocation method of a kernel defined as pointer to a kernel object, described in detail in 3.5.3
<pre>void parallel_for(range<dimensions> numWorkItems, kernel sycl_kernel)</pre>	Kernel invocation method of a kernel defined as pointer to a kernel object, for the specified range and given an id or item for indexing in the indexing space defined by range, described in detail in 3.5.3
<pre>void parallel_for(nd_range<dimensions> ndRange, kernel syclKernel)</pre>	Kernel invocation method of a kernel defined as pointer to a kernel object, for the specified <code>nd_range</code> and given an <code>nd_item</code> for indexing in the indexing space defined by the <code>nd_range</code> , described in detail in 3.5.3
End of table	

Table 3.59: Methods for the command group `handler` class

-In the case of OpenCL/SYCL interoperability the `handler` class provides a method for setting kernel arguments for a kernel that is built from the OpenCL host interface or other libraries, like CLU. The developer may have already set the kernel arguments or may choose to use this method, keeping in mind that this method is not thread safe, as per the OpenCL specification.

3.5.4 Kernel class

The *kernel class* is an abstraction of a kernel object in SYCL. At 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 3.5.3. 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 program class, as defined in 3.5.5. In both the above cases, the developer cannot instantiate a kernel object but can instantiate an object a functor class that he could use or create a functor from a kernel method using C++11 features. The kernel class object needs a `parallel_for(...)` invocation or an explicit `compile_and_link()` call through the program class, for this compilation of the kernel to be triggered.

Finally, a kernel class instance may encapsulate an OpenCL *kernel* object that was created using the OpenCL C interface and the arguments of the kernel already set. In this case since the developer is providing the `cl_kernel` object, this constructor is allowed to be used by the developer.

The kernel class also provides the interface for getting information from a kernel object. The kernel information descriptor interface is described in C.5 and the description is in the table 3.62.

```
1 namespace cl {
2 namespace sycl {
3 class kernel {
4 private:
5     friend class program;
6
7     // The default object is not valid because there is no
8     // program or cl_kernel associated with it
9     kernel();
10
11 public:
12     kernel(const kernel& rhs);
13
14     kernel(cl_kernel openclKernelObject);
15
16     cl_kernel get() const;
17
18     context get_context() const;
19
20     program get_program() const;
21
22     template <info::kernel param>
23     typename info::param_traits<info::kernel, param>::type
24     get_info() const;
25 };
26 } // namespace sycl
27 } // namespace cl
```

Constructor	Description
<code>kernel (cl_kernel openclKernelObj)</code>	Constructor for SYCL kernel class given an OpenCL kernel object with set arguments, valid for enqueueing. Retains a reference to the <code>cl_kernel</code> object. Caller should release the passed <code>cl_kernel</code> object when it is no longer needed.
<code>kernel (const kernel& rhs)</code>	Copy constructor for kernel class.
End of table	

Table 3.60: `kernel` class constructors

Methods	Description
<code>cl_kernel get()</code>	Return the OpenCL kernel object for this kernel. Retains a reference to the returned <code>cl_kernel</code> object. Caller should release it when finished.
<code>context get_context()</code>	Return the context that this kernel is defined for.
<code>program get_program()</code>	Return the program that this kernel is part of.
<code>template <info::kernel param> typename info::param_traits< info::kernel, param>::type get_info()const</code>	Query information from the kernel object using the <code>info::kernel_info</code> descriptor.
End of table	

Table 3.61: Methods for the `kernel` class.

Kernel Descriptors	Return type	Description
<code>info::kernel::function_name</code>	<code>string_class</code>	Return the kernel function name.
<code>info::kernel::num_args</code>	<code>cl_uint</code>	Return the number of arguments to the extracted OpenCL C kernel.
<code>info::kernel::reference_count</code>	<code>cl_uint</code>	Return the reference count of the kernel object.
<code>info::kernel::attributes</code>	<code>string_class</code>	Return any attributes specified using the <code>__attribute__</code> qualifier with the kernel function declaration in the program source.
End of table		

Table 3.62: Kernel class information descriptors.

3.5.5 Program class

A *program* contains one or more kernels and any functions or libraries necessary for the program's execution. A program will be enqueued inside a context and each of the kernels will be enqueued on a corresponding device. Program class can be really useful for pre-compiling kernels and enqueueing them on multiple command groups. It also allows usage of functions define in OpenCL kernels from SYCL kernels via compiling and linking them in

the same program object.

The program class provides an interface for getting information from a program object. The program information descriptor interface is described in [C.6](#) and the description is in the table [3.65](#).

```
1 namespace cl {
2 namespace sycl {
3 class program {
4 public:
5 // Create an empty program object
6 explicit program(const context& context);
7
8 // Create an empty program object
9 program(const context& context, vector_class<device> deviceList);
10
11 // Create a program object from a cl_program object
12 program(const context& context, cl_program clProgram);
13
14 // Create a program by linking a list of other programs
15 program(vector_class<program> programList, string_class linkOptions = "");
16
17 program(const program& rhs);
18
19 ~program();
20
21 /* This obtains a SYCL program object from a SYCL kernel name
22 and compiles it ready to link */
23 template <typename kernelT>
24 void compile_from_kernel_name(string_class compileOptions = "");
25
26 /* This obtains a SYCL program object from a SYCL kernel name
27 and builds it ready-to-run */
28 template <typename kernelT>
29 void build_from_kernel_name(string_class compileOptions = "");
30
31 void link(string_class linkingOptions = "");
32
33 // Get a kernel from a given Name (Functor)
34 template <typename kernelT>
35 kernel get_kernel<kernelT>() const;
36
37 template <info::program param>
38 typename info::param_traits<info::program, param>::type
39 get_info() const;
40
41 vector_class<vector_class<char>> get_binaries() const;
42
43 vector_class<::size_t> get_binary_sizes() const;
44
45 vector_class<device> get_devices() const;
46
47 string_class get_build_options() const;
48
49 cl_program get() const;
50
```

```

51  bool is_linked() const;
52  };
53  } // namespace sycl
54  } // namespace cl

```

Constructors	Description
<code>explicit program (const context & context)</code>	Constructs an empty program object for context for all associated devices with context.
<code>program (const context & context, vector_class<device> deviceList)</code>	Constructs an empty program object for all the devices of <i>device_list</i> associated with the <i>context</i> .
<code>program (vector_class<program> programList, string_class linkOptions="")</code>	Constructs a program object for a list of programs and links them together using the <i>linkOptions</i> .
<code>program (const context & context, cl_program clProgram)</code>	Constructs a program object for an OpenCL program object. Retains a reference to the returned <i>cl_program</i> object. Calling context should release it when finished.
<code>program(const program& rhs);</code>	Copy constructor for the program class.
End of table	

Table 3.63: Constructors for the `program` class

Methods	Description
<code>template<typename kernelT> void compile_from_kernel_name(string_class compileOptions="")</code>	Compile the kernel defined to be of type <i>kernelT</i> into the program, with compile options given by <i>compile_options</i> ". The kernel can be defined either as a functor of type <i>kernelT</i> or as a lambda function which is named with the class name <i>kernelT</i> . The program object will need to be linked later.
<code>template<typename kernelT> void build_from_kernel_name(string_class compileOptions="")</code>	Build the kernel defined to be of type <i>kernelT</i> into the program, with compile options given by <i>compile_options</i> ". The kernel can be defined either as a functor of type <i>kernelT</i> or as a lambda function which is named with the class name <i>kernelT</i> .
<code>void link(string_class linking_options = "")</code>	Link all compiled programs that are added in the program class.
<code>template<info::program param> typename info::param_traits<info::program, param>::type get_info()</code>	Retrieve information of the built OpenCL program object.
<code>vector_class<char*> get_binaries() const</code>	Return the array of compiled binaries associated with the program, as compiled for each device.
<code>vector_class<device> get_devices() const</code>	Return the list of devices this program was constructed against
Continued on next page	

Table 3.64: Methods for the `program` class

Methods	Description
<code>string_class</code> <code>get_build_options()const</code>	Retrieve the set of build options of the program. A program is created with one set of build options.
<code>cl_program</code> <code>get()const</code>	Return the OpenCL program object for this program. Retains a reference to the returned <code>cl_program</code> object. Caller should release it when finished.
End of table	

Table 3.64: Methods for the `program` class

Program Descriptors	Return type	Description
<code>info::program::reference_count</code>	<code>cl_uint</code>	Return the reference count of the kernel object.
<code>info::program::context</code>	<code>cl_context</code>	Return the context object this program is associated with.
<code>info::program::devices</code>	<code>vector_class<cl_device_id></code>	Return set of devices this program is built against.
End of table		

Table 3.65: Program class information descriptors.

Programs allow the developers to provide their own compilation and linking options and also compile and link on demand one or multiple kernels. Compiler options allowed are described in the OpenCL specification [1, p. 145, § 5.6.4] and the linker options are described in [1, p. 148, § 5.6.5].

3.5.6 Defining kernels

In SYCL functions that are executed in parallel on a SYCL device are referred to as *kernel functions*. A *kernel* containing such a *kernel function* is enqueued on a device queue in order to be executed on that particular device. The return type of the *kernel function* is `void`, and all kernel accesses between host and device are defined using the accessor class 3.4.6.

There are three ways of defining kernels, defining them as functors, as C++11 lambda functions or as OpenCL `cl_kernel` objects. However, in the case of OpenCL kernels, the developer is expected to have created the kernel and set the kernel arguments.

3.5.6.1 Defining kernels as functors

A kernel can be defined as a C++ functor. In this case, the *kernel function* is the function defined as `operator()` in the normal C++ functor style. These functors provide the same functionality as any C++ functors, with the restriction that they need to follow C++11 standard layout rules. The kernel function can be templated via templating the kernel functor class. The `operator()` function may take different parameters depending on the data accesses that defined for the specific kernel. For details on restrictions for kernel naming issues, please refer to 5.2.

In the following example we define a trivial functor with no outputs and no accesses of host or pre-allocated

device data. The kernel is executed on a unary index space for the specific example, since its using `single_task` at its invocation.

```
1 class MyFunctor
2 {
3     float m_parameter;
4
5 public:
6     MyFunctor(float parameter):
7         m_parameter(parameter)
8     {
9     }
10
11     void operator() ()
12     {
13         // [kernel code]
14     }
15 };
16
17 void workFunction(float scalarValue)
18 {
19     MyFunctor myKernel(scalarValue);
20
21     queue.submit([& (handler & cmdGroup) {
22         cmdgroup.single_task(myKernel);
23     }]);
24 }
```

3.5.6.2 Defining kernels as lambda functions

In C++11, functors can be defined using lambda functions. We allow lambda functions to define kernels in SYCL, but we have an extra requirement to *name lambda functions* in order to enable the linking of the SYCL device kernels with the host code to invoke them. The name of a lambda function in SYCL is a C++ class. If the lambda function relies on template arguments, then 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 issues, please refer to [5.2](#).

To invoke a C++11 lambda, the kernel name must be included explicitly by the user as a template parameter to the kernel invoke function.

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. [=]).

```
1 class MyKernel;
2
3 command_queue.submit([& (handler& cmdGroup) {
4     cmdgroup.single_task<class MyKernel>([=]) () {
5         // [kernel code]
6     }]);
7 }]);
```

3.5.6.3 Defining kernels using program objects

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 3.5.5. The kernel is defined as a functor 3.5.6.1 or lambda function 3.5.6.2. The user can obtain a program object for the kernel with the `get_kernel` method. This method is templated by the *kernel name*, so that the user can specify the kernel whose associated kernel they wish to obtain.

In the following example, the kernel is defined as a lambda function. The example obtains the program object for the lambda function kernel and then passes it to the `parallel_for`.

```
1 class MyKernel; // Forward declaration of the name of the lambda functor
2
3 cl::sycl::queue myQueue;
4 cl::sycl::program MyProgram(myQueue.get_context());
5
6 /* use the name of the kernel to obtain the associated program */
7 MyProgram.build_from_name<MyKernel>();
8
9 myQueue.submit([& (handler& commandGroup) {
10     commandGroup.parallel_for<class MyKernel>(
11         cl::sycl::nd_range<2>(4, 4),
12         MyProgram.get_kernel<MyKernel>(), // execute the kernel as compiled in MyProgram
13         ( [=](cl::sycl::nd_item<2> index) {
14             //[kernel code]
15         }));
16 });
```

In the above example, the *kernel function* is defined in the `parallel_for` invocation as part of a lambda functor which is named using the type of the forward declared class “myKernel”. The type of the functor and the program object enable the compilation and linking of the kernel in the program class, *a priori* of its actual invocation as a kernel object. For more details on the SYCL device compiler please refer to chapter 5.

In the next example, a SYCL kernel is linked with an existing pre-compiled OpenCL C program object to create a combined program object, which is then called in a `parallel_for`.

```
1 class MyKernel; // Forward declaration of the name of the lambda functor
2
3 cl::sycl::queue myQueue;
4
5 // obtain an existing OpenCL C program object
6 cl_program myClProgram = ...;
7
8 // Create a SYCL program object from a cl_program object
9 cl::sycl::program myExternProgram(myQueue.get_context(), myClProgram);
10
11 // Release the program if we no longer need it as
12 // SYCL program retained a reference to it
13 clReleaseProgram(myClProgram);
14
15 // Add in the SYCL program object for our kernel
16 cl::sycl::program mySyclProgram(myQueue.get_context());
17 mySyclProgram.compile_from_kernel_name<MyKernel>("-my-compile-options");
```

```

18
19 // Link myClProgram with the SYCL program object
20 mySyclProgram.link(myExternProgram, "-my-link-options");
21
22 myQueue.submit([&](handler& commandgroup) {
23     commandgroup.parallel_for<class MyKernel>(
24         cl::sycl::nd_range<2>(4, 4),
25         myLinkedProgram.get_kernel<MyKernel>(), // execute the kernel as compiled in MyProgram
26         ([=](cl::sycl::nd_item<2> index) {
27             //[kernel code]
28         }));
29 });

```

3.5.6.4 Defining kernels using OpenCL C kernel objects

In OpenCL C [1] program and kernel objects can be created using the OpenCL C API, which is available in the SYCL system. Interoperability of OpenCL C kernels and the SYCL system is achieved by allowing the creation of a *SYCL kernel* object from an *OpenCL kernel* object.

The constructor using kernel objects from 3.60:

```
kernel::kernel(cl_kernel kernel)
```

creates a `cl::sycl::kernel` which can be enqueued using all the `parallel_for` functions which can enqueue a kernel object. This way of defining kernels assumes the developer is using OpenCL C to create the kernel and set the kernel arguments. The system assumes that the developer has already called set kernel arguments when they are trying to enqueue the kernel. Buffers do give ownership to their accessors on specific contexts and the developer can enqueue OpenCL kernels in the same way as enqueueing SYCL kernels. However, the system is not responsible for data management at this point. Note that like all constructors from OpenCL API objects, constructing a `cl::sycl::kernel` from a `cl_kernel` will retain a reference to the kernel and the user code should call `clReleaseKernel` if the `cl_kernel` is no longer needed in the calling context.

3.5.7 Rules for parameter passing to kernels

In a case where a kernel is a C++ functor or C++11 lambda object, any values in the functor or captured in the C++11 lambda object 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. For OpenCL 1.0–1.2 class devices, this means that the argument must be passed via `clSetKernelArg` and be compiled as a kernel parameter of the valid reference type. For global shared data access, the parameter must be an OpenCL `global` pointer. For an accessor that specifies OpenCL constant access, the parameter must be an OpenCL constant pointer. For images, the accessor must be passed as an `image_t` and/or sampler.
- 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.
- 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++ standard layout values must be passed by value to the kernel.
- C++ non-standard layout values 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 (`cl::sycl::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 OpenCL kernel parameters if all members of the aggregate type are unaffected by the rules above.

3.6 Error handling

3.6.1 Error Handling Rules

Error handling in SYCL uses exceptions. If an error can be propagated at the point of calling a function, an exception will be thrown and may be caught by the user using standard C++ exception handling mechanisms. For example, any exception which is triggered from code executed on host is able to be propagated at the call site and it will follow the standard C++ exception handling mechanisms.

SYCL applications are asynchronous in the sense that host and device code execution are executed asynchronously. As a result of this, the errors that occur on a device cannot be propagated directly from the call site, and they will not be detected until the error-causing task executes or tries to execute rather than been scheduled for execution. We refer to those errors as asynchronous errors. A good example of an asynchronous error, is an out-of-bounds access error. In this case, if the kernel is enqueued on SYCL OpenCL device then the out-of-bounds error is asynchronous with regards to the SYCL host application, as it is executed on the device. At the latter, the standard exception mechanisms will not be available as this is an asynchronous error.

SYCL queues are by default asynchronous, as they schedule tasks on SYCL devices. The queue constructor can optionally get an asynchronous handler object, which can be a functor or lambda function for receiving a list of C++ exception objects when the waiting and exception handling methods are called on queues.

If an asynchronous error occurs in a queue that has no user-supplied asynchronous error handler object, then no exception is thrown and the error is not available to the user in any specified way. Implementations may provide extra debugging information to users to trap and handle asynchronous errors. If a synchronous error occurs in a SYCL application and it is not handled, the application will exit abnormally.

If an error occurs when running or en-queuing 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 asynchronous handler object provided for both queues. If there is no asynchronous handler object 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 asynchronous handler object 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 asynchronous error handler object and either `wait_and_throw()` or `throw()` are called on that queue. The command group functor handler event returned by that function will be relevant to the queue where the kernel has been enqueued.

3.6.2 Exception Class Interface

```

typedef function_class<void(cl::sycl::exception_list)> async_handler;

namespace cl {
namespace sycl {

struct exception {
    string_class what();

    // returns associated context. nullptr if none
    context get_context();
};

struct cl_exception : exception {
    // thrown as a result of an OpenCL API error code
    cl_int get_cl_code() const;
};

struct async_exception : exception {
    // stored in an exception_list for asynchronous errors
};

class exception_list {
    // Used as a container for a list of asynchronous exceptions
public:
    typedef exception_ptr value_type;
    typedef value_type& reference;
    typedef const value_type& const_reference;
    typedef size_t size_type;
    typedef /*unspecified*/ iterator;
    typedef /*unspecified*/ const_iterator;

    size_t size() const;
    iterator begin() const; // first asynchronous exception
    iterator end() const; // refer to past-the-end last asynchronous exception
};

typedef /*unspecified*/ exception_ptr;

function_class<void(cl::sycl::exception_list)> async_handler;

class runtime_error : public exception;

```

```

class kernel_error : public runtime_error;

class accessor_error: public runtime_error;

class nd_range_error : public runtime_error;

class event_error : public runtime_error;

class invalid_parameter_error : public runtime_error;

class device_error : public exception;

class compile_program_error : public device_error;

class link_program_error : public device_error;

class invalid_object_error : public device_error;

class memory_allocation_error : public device_error;

class platform_error : public device_error;

class profiling_error : public device_error;

class feature_not_supported : public device_error;

} // namespace sycl
} // namespace cl

```

The `cl_exception` class is the exception thrown when the OpenCL API returns an error code. The OpenCL error code can be queried with the `get_cl_code` method. The `async_exception` is stored in `exception_list` objects and is generated when an asynchronous error occurs on a SYCL-managed context. The `cl::sycl::exception_ptr` class is used to store `cl::sycl::exception` objects and allows exception objects to be transferred between threads. It is equivalent to the `std::exception_ptr` class. The `cl::sycl::exception_list` class is also available.

The asynchronous handler object, `async_handler`, can be a `function_class` with an `exception_list` as a parameter. The asynchronous handler is an optional parameter to a constructor of the `queue` class and it is the only way to handle asynchronous errors occurring on a SYCL device. The asynchronous handler may be a `function_class` that can be functor or lambda or function that can be given to the queue and it will be executed on error. The `exception_list` object is constructed from the SYCL runtime and is populated with the errors caught during the execution of all the kernels running on the same queue.

Methods	Description
<code>string_class</code> <code>what()</code>	Returns a descriptive string for the error, if available.
<code>context</code> <code>get_context()</code>	Returns the context that caused the error. Returns <code>nullptr</code> if not a buffer error.
End of table	

Table 3.66: Methods of the `exception` class.

Methods	Description
<code>cl_int get_cl_code()</code>	Returns the OpenCL error code encapsulated in the exception. Only valid for the <code>cl_exception</code> subclass.
End of table	

Table 3.67: Methods of the `cl_exception` class.

Methods	Description
<code>size_t size()const</code>	Returns the size of the list
<code>iterator begin()const</code>	Returns an iterator to the beginning of the list of asynchronous exceptions.
<code>iterator end()const</code>	Returns an iterator to the end of the list of asynchronous exceptions.
End of table	

Table 3.68: Methods of the `exception_list`

Runtime Error Exception Type	Description
<code>kernel_error</code>	Error that occurred before or while enqueueing the SYCL kernel.
<code>nd_range_error</code>	Error regarding the <code>cl::sycl::nd_range</code> specified for the SYCL kernel
<code>accessor_error</code>	Error regarding the <code>cl::sycl::accessor</code> objects defined.
<code>event_error</code>	Error regarding associated <code>cl::sycl::event</code> objects.
<code>invalid_parameter_error</code>	Error regarding parameters to the SYCL kernel, it may apply to any captured parameters to the kernel lambda.
End of table	

Table 3.69: Exceptions types that derive from the `cl::sycl::runtime_error` class

Device Error Exception Type	Description
<code>compile_program_error</code>	Error while compiling the SYCL kernel to a SYCL device.
<code>link_program_error</code>	Error while linking the SYCL kernel to a SYCL device.
<code>invalid_object_error</code>	Error regarding any memory objects being used inside the kernel
<code>memory_allocation_error</code>	Error on memory allocation on the SYCL device for a SYCL kernel.
<code>device_error</code>	The SYCL device will trigger this exception on error.
Continued on next page	

Table 3.70: Exception types that derive from the `cl::sycl::device_error` class

Device Error Exception Type	Description
platform_error	The SYCL platform will trigger this exception on error.
profiling_error	The SYCL runtime will trigger this error if there is an error when profiling info is enabled.
feature_not_supported	Exception thrown when an optional feature or extension is used in a kernel but its not available on the device the SYCL kernel is being enqueued on.
End of table	

Table 3.70: Exception types that derive from the `cl::sycl::device_error` class

3.7 Data types

SYCL as a C++11 programming model supports the C++11 ISO standard data types, and it also provides the ability for all SYCL applications to be executed on SYCL compatible devices, OpenCL and host 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](#)

3.7.1 Scalar data types

SYCL follows the C++11 standard in terms of fundamental scalar data types. All SYCL applications match those data types and the size of those has to be matched in the SYCL application code for all available SYCL devices; host and OpenCL devices.

SYCL Integral Data Types	Description
<code>char</code>	a signed 8-bit integer, as defined by the C++11 ISO Standard
<code>unsigned char</code>	an unsigned 8-bit integer, as defined by the C++11 ISO Standard
<code>short int</code>	a signed integer of at least 16-bits, as defined by the C++11 ISO Standard
<code>unsigned short int</code>	an unsigned integer of at least 16-bits, as defined by the C++11 ISO Standard
<code>int</code>	a signed integer of at least 16-bits, as defined by the C++11 ISO Standard
<code>unsigned int</code>	an unsigned integer of at least 16-bits, as defined by the C++11 ISO Standard
<code>long int</code>	a signed integer of at least 32-bits, as defined by the C++11 ISO Standard
<code>unsigned long int</code>	an unsigned integer of at least 32-bits, as defined by the C++11 ISO Standard
Continued on next page	

Table 3.71: SYCL compiler fundamental integral datatypes

SYCL Integral Data Types	Description
<code>long long int</code>	an integer of at least 64-bits, as defined by the C++11 ISO Standard
<code>unsigned long long int</code>	an unsigned integer of at least 64-bits, as defined by the C++11 ISO Standard
<code>size_t</code>	the unsigned integer type of the result of the <code>sizeof</code> operator on host.
End of table	

Table 3.71: SYCL compiler fundamental integral datatypes

SYCL Floating Point Data Types	Description
<code>float</code>	a 32-bit IEEE 754 floating-point value, as defined by the C++11 ISO Standard
<code>double</code>	a 64-bit IEEE 754 floating-point value, as defined by the C++11 ISO Standard
<code>half</code>	a 16-bit IEEE 754-2008 half-precision floating-point value
End of table	

Table 3.72: SYCL compiler fundamental floating point datatypes

The OpenCL C language standard [par. 6.11] `opencl-1.2` defines its own built-in scalar types, which are supported for interoperability between SYCL and OpenCL C applications through the interoperability data types.

SYCL Scalar Datatypes	Description
<code>cl::sycl::cl_bool</code>	A conditional data type which is either true or false. The value true expands to the integer constant 1 and the value false expands to the integer constant 0.
<code>cl::sycl::cl_char</code>	a signed two's complement 8-bit integer
<code>cl::sycl::cl_uchar</code>	an unsigned 8-bit integer
<code>cl::sycl::cl_short</code>	a signed two's complement 16-bit integer
<code>cl::sycl::cl_ushort</code>	an unsigned 16-bit integer
<code>cl::sycl::cl_int</code>	a signed two's complement 32-bit integer
<code>cl::sycl::cl_uint</code>	an unsigned 32-bit integer
<code>cl::sycl::cl_long</code>	A signed two's complement 64-bit integer.
<code>cl::sycl::cl_ulong</code>	An unsigned 64-bit integer.
<code>cl::sycl::cl_float</code>	A 32-bit floating-point. The float data type must conform to the IEEE 754 single precision storage format.
<code>cl::sycl::cl_double</code>	A 64-bit floating-point. The double data type must conform to the IEEE 754 double precision storage format.
<code>cl::sycl::cl_half</code>	A 16-bit floating-point. The half data type must conform to the IEEE 754-2008 half precision storage format.
End of table	

Table 3.73: SYCL compiler OpenCL interoperability scalar datatypes

3.7.2 Vector types

SYCL provides a templated cross-platform vector type 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 methods that allow construction of a new vector from a swizzled set of component elements. The vector are defined in Table 3.75 and Table 3.76

An overview of the interface provided for the `vec` class is the following, for the full interface please refer to B.5.

```
namespace cl {
namespace sycl {
template <typename dataT, int numElements>
class vec {
public:
    typedef dataT element_type;
    //Underlying OpenCL type
    typedef __undefined__ vector_t;

    vec();

    explicit vec(const dataT &arg);

    vec(const T0 &arg0... args);

    vec(const vec<dataT, numElements> &rhs);

    size_t get_count();

    size_t get_size();

    template <typename asDataT, int width>
    vec<asDataT, width> as() const;

    // genvector is a generic typename for describing
    // all OpenCL/SYCL types.
    operator __genvector__() const;

    // Swizzle methods (see notes)
    swizzled_vec<T, out_dims> swizzle<elem s1, ...>();

#ifdef SYCL_SIMPLE_SWIZZLES
    swizzled_vec<T, 4> xyzw();
    ...
#endif // #ifdef SYCL_SIMPLE_SWIZZLES
};
} // namespace sycl
} // namespace cl
```

`vec<typename T, int dims>` is a vector type that compiles down to the OpenCL built-in vector types on OpenCL devices where possible and provides compatible support on the host. The `vec` class is templated on its number of dimensions and its element type. The dimensions parameter, `dims`, can be one of: 1, 2, 3, 4, 8 or 16. Any other value should produce a compilation failure. The element type parameter, `T`, must be one of the basic scalar types supported in device code.

The SYCL library provides typedefs for: `char`, `unsigned char`, `short`, `unsigned short`, `int`, `unsigned int`, `long`, `long long`, `unsigned long long`, `float` and `double` in all valid sizes. These vector typedefs are named `TypenameSize`, for example: `longlong2` is a vector of two long long integer elements, mapping to `vec<long long int, 2>`.

The SYCL library also provides the OpenCL interoperability types which are referred to in the document with the generic name `genvector`. The following table 3.74 is showing the the types that are available and are represented for brevity in the specification as a generic and not an actual SYCL type `genvector`.

Generic type name	Description
<code>genvector</code>	<code>cl::sycl::cl_float2, cl::sycl::cl_float3, cl::sycl::cl_float4, cl::sycl::cl_float8, cl::sycl::cl_float16</code> <code>cl::sycl::cl_double2, cl::sycl::cl_double3, cl::sycl::cl_double4, cl::sycl::cl_double8, cl::sycl::cl_double16</code> <code>cl::sycl::cl_char2, cl::sycl::cl_char3, cl::sycl::cl_char4, cl::sycl::cl_char8, cl::sycl::cl_char16</code> <code>cl::sycl::cl_uchar2, cl::sycl::cl_uchar3, cl::sycl::cl_uchar4, cl::sycl::cl_uchar8, cl::sycl::cl_uchar16</code> <code>cl::sycl::cl_short2, cl::sycl::cl_short3, cl::sycl::cl_short4, cl::sycl::cl_short8, cl::sycl::cl_short16</code> <code>cl::sycl::cl_ushort2, cl::sycl::cl_ushort3, cl::sycl::cl_ushort4, cl::sycl::cl_ushort8, cl::sycl::cl_ushort16</code> <code>cl::sycl::cl_uint2, cl::sycl::cl_uint3, cl::sycl::cl_uint4, cl::sycl::cl_uint8, cl::sycl::cl_uint16</code> <code>cl::sycl::cl_int2, cl::sycl::cl_int3, cl::sycl::cl_int4, cl::sycl::cl_int8, cl::sycl::cl_int16</code> <code>cl::sycl::cl_ulong2, cl::sycl::cl_ulong3, cl::sycl::cl_ulong4, cl::sycl::cl_ulong8, cl::sycl::cl_ulong16</code> <code>cl::sycl::cl_long2, cl::sycl::cl_long3, cl::sycl::cl_long4, cl::sycl::cl_long8, cl::sycl::cl_long16</code>
End of table	

Table 3.74: Generic type name description for `genVector`, which serves as a description for all valid types of OpenCL/SYCL interoperability vectors.

`swizzled_vec<T, out_dims> vec<T, in_dims>::swizzle<elem s1, elem s2...> ()` returns a temporary object representing a swizzled set of the original vector's member elements. The number of `s1, s2` param-

eters is the same as `out_dims`. All `s1`, `s2` parameters must be integer constants from zero to `in_dims-1`. The swizzled vector may be used as a source (r-value) and destination (l-value). In order to enable the r-value and l-value swizzling to work, this returns an intermediate swizzled-vector class, which can be implicitly converted to a vector (r-value evaluation) or assigned to.

If the user `#defines` the macro `SYCL_SIMPLE_SWIZZLES` before `#include <cl/sycl.hpp>`, then swizzle functions are defined for every combination of swizzles for 2D, 3D and 4D vectors only. The swizzle functions take the form:

```
swizzled_vec<T, out_dims> vec<T, in_dims>::xyzw();
swizzled_vec<T, out_dims> vec<T, in_dims>::rgba();
```

where, as above, the number of `elem::x`, `elem::y`, `elem::z`, `elem::w` or `elem::r`, `elem::g`, `elem::b`, `elem::a` letters is the same as `out_dims`. All `elem::x`, `elem::y`, `elem::z`, `elem::w` or `elem::r`, `elem::g`, `elem::b`, `elem::a` parameters must be letters from the sets first `in_dims` letters in “xyzw” or “rgba”.

Swizzle letters may be repeated or re-ordered. For example, from a vector contains integers [0, 1, 2, 3], `vec.xxzy()` would return a vector of [0, 0, 2, 1].

Constructors	Description
<code>vec<T, dims>()</code>	Default construct a vector with element type T and with <code>dims</code> dimensions by default construction of each of its elements.
<code>explicit vec<T, dims>(const T &arg)</code>	Construct a vector of element type T and <code>dims</code> dimensions by setting each value to <code>arg</code> by assignment.
<code>vec<T, dims> (</code> <code> const T &element_0,</code> <code> const T &element_1,</code> <code> . . . ,</code> <code> const T &element_dims-1)</code>	Construct a vector with element type T and with <code>dims</code> dimensions out of <code>dims</code> initial values.
<code>vec<T, dims>(const &vec<T, dims>)</code>	Construct a vector of element type T and <code>dims</code> dimensions by copy from another similar vector.
End of table	

Table 3.75: Constructors for the `vec` class

Methods	Description
<code>size_t get_count()</code>	Returns the number of elements of the vector.
<code>size_t get_size()</code>	Returns the size of the vector.
<code>template<typename asDataT, int width></code> <code>vec<asDataT,width> as()const</code>	Re-interpret the <code>vec<dataT,numElements></code> to <code>vec<asDataT,width></code> type.
Continued on next page	

Table 3.76: Methods for the `vec` class

Methods	Description
<code>operator genvector() const</code>	Converts a <code>vec<dataT, numElements></code> to the corresponding OpenCL vector, generically named <code>openclVector</code> ¹ of the same type and width.
<code>operator vec<dataT, numElements>(genvector clVector) const</code>	Assignment operator that takes an OpenCL vector instance, which one of <code>openclVector</code> , and converts it to the corresponding SYCL <code>vec<dataT, numElements></code> type.
<code>vec<dataT, numElements> operator+(const vec<dataT, numElements> &rhs)</code>	Construct vector from the sum of the respective elements of the current vector and rhs.
<code>vec<dataT, numElements> operator+(const dataT &rhs)</code>	Construct vector by adding rhs to each element of the current vector.
<code>vec<dataT, numElements> operator-(const vec<dataT, numElements> &rhs)</code>	Construct vector by subtracting the elements of rhs from the respective elements of the current vector.
<code>vec<dataT, numElements> operator-(const dataT &rhs)</code>	Construct vector by subtracting rhs from each element of the current vector.
<code>vec<dataT, numElements> operator*(const vec<dataT, numElements> &rhs)</code>	Construct vector from the product of the elements of the current vector by the respective elements of rhs.
<code>vec<dataT, numElements> operator*(const dataT &rhs)</code>	Construct vector by multiplying each element of the current vector by rhs.
<code>vec<dataT, numElements> operator/(const vec<dataT, numElements> &rhs)</code>	Construct vector from the division of the elements of the current vector by the elements of rhs.
<code>vec<dataT, numElements> operator/(const dataT &rhs)</code>	Construct vector by dividing each element of the current vector by rhs.
<code>vec<dataT, numElements> operator%(const vec<dataT, numElements> &rhs)</code>	Construct vector from the modulo of the elements of the current vector and the elements of rhs.
<code>vec<dataT, numElements> operator%(const dataT &rhs)</code>	Construct vector by calculating the remainder of each element of the current vector and rhs.
<code>vec<dataT, numElements> operator++();</code>	Prefix increment by one for every element of the vector.
<code>vec<dataT, numElements> operator++(int);</code>	Post-fix increment by one for every element of the vector.
<code>vec<dataT, numElements> operator--();</code>	Prefix decrement by one for every element of the vector.
<code>vec<dataT, numElements> operator--(int);</code>	Post-fix decrement by one for every element of the vector.
<code>vec<dataT, numElements> operator (const vec<dataT, numElements> &rhs)</code>	Construct vector from the bitwise OR of the respective elements of the current vector and rhs.
<code>vec<dataT, numElements> operator (const dataT &rhs)</code>	Construct vector by applying the bitwise OR of each element of the current vector and rhs.

Continued on next page

Table 3.76: Methods for the `vec` class

¹This is not actual SYCL type, the `genvector` type is described in table 3.74

Methods	Description
<code>vec<dataT, numElements> operator^(const vec<dataT, numElements> &rhs)</code>	Construct vector from the bitwise AND of the respective elements of the current vector and rhs.
<code>vec<dataT, numElements> operator^(const dataT &rhs)</code>	Construct vector by applying the bitwise AND of each element of the current vector and rhs.
<code>vec<dataT, numElements> operator&&(</code> <code>const vec<dataT, numElements> &rhs)</code>	Construct vector from the logical AND of the respective elements of the current vector and rhs.
<code>vec<dataT, numElements> operator&&(</code> <code>const dataT &rhs)</code>	Construct vector by applying the logical AND of each element of the current vector and rhs.
<code>vec<dataT, numElements> operator (</code> <code>const vec<dataT, numElements> &rhs)</code>	Construct vector from the logical OR of the respective elements of the current vector and rhs.
<code>vec<dataT, numElements> operator (</code> <code>const dataT &rhs)</code>	Construct vector by applying the logical OR of each element of the current vector and rhs.
<code>vec<dataT, numElements> operator>>(</code> <code>const vec<dataT, numElements> &rhs)</code>	Construct vector from the outcome of shifting right the respective elements of the current vector by rhs.
<code>vec<dataT, numElements> operator>>(</code> <code>const dataT &rhs)</code>	Construct vector by shifting right each element of the current vector by rhs.
<code>vec<dataT, numElements> operator<<(</code> <code>const vec<dataT, numElements> &rhs)</code>	Construct vector from the outcome of shifting left the respective elements of the current vector by rhs.
<code>vec<dataT, numElements> operator<<(</code> <code>const dataT &rhs)</code>	Construct vector by shifting left each element of the current vector by rhs.
<code>vec<dataT, numElements> operator~()</code>	Construct vector from the outcome of applying bitwise not to the respective elements of the current vector.
<code>vec<dataT, numElements> operator!()</code>	Construct vector from the outcome of applying logical not to the respective elements of the current vector.
<code>vec<dataT, numElements> operator+=(</code> <code>const vec<dataT, numElements> &rhs)</code>	Add each element of rhs to the respective element of the current vector in-place.
<code>vec<dataT, numElements> operator+=(</code> <code>const dataT &rhs)</code>	Add rhs in-place to each element of the current vector.
<code>vec<dataT, numElements> operator-=(</code> <code>const vec<dataT, numElements> &rhs)</code>	Subtract each element of rhs from the respective element of the current vector in-place.
<code>vec<dataT, numElements> operator-=(</code> <code>const dataT &rhs)</code>	Subtract rhs in-place from each element of the current vector.
<code>vec<dataT, numElements> operator*=(</code> <code>const vec<dataT, numElements> &rhs)</code>	Multiple each element of the current vector by the respective element of rhs in-place.
<code>vec<dataT, numElements> operator*=(</code> <code>const dataT &rhs)</code>	Multiple in-place each element of the current vector by rhs.

Continued on next page

Table 3.76: Methods for the vec class

Methods	Description
<code>vec<dataT, numElements> operator/=(const vec<dataT, numElements> &rhs)</code>	Divide each element of the current vector in-place by the respective element of rhs.
<code>vec<dataT, numElements> operator/=(const dataT &rhs)</code>	Divide in-place each element of the current vector by rhs.
<code>vec<dataT, numElements> operator =(const vec<dataT, numElements> &rhs)</code>	Bitwise OR of each element of the current vector in-place by the respective element of rhs.
<code>vec<dataT, numElements> operator =(const dataT &rhs)</code>	Bitwise OR in-place each element of the current vector by rhs.
<code>vec<dataT, numElements> operator^=(const vec<dataT, numElements> &rhs)</code>	Bitwise and of each element of the current vector in-place by the respective element of rhs.
<code>vec<dataT, numElements> operator^=(const dataT &rhs)</code>	Bitwise AND in-place each element of the current vector by rhs.
<code>vec<dataT, numElements> operator<<=(const vec<dataT, numElements> &rhs)</code>	Shift left each element of the current vector in-place by the respective element of rhs.
<code>vec<dataT, numElements> operator<<=(const dataT &rhs)</code>	Shift left in-place each element of the current vector by rhs.
<code>vec<dataT, numElements> operator>>=(const vec<dataT, numElements> &rhs)</code>	Shift right each element of the current vector in-place by the respective element of rhs.
<code>vec<dataT, numElements> operator>>=(const dataT &rhs)</code>	Shift right in-place each element of the current vector by rhs.
<code>vec<dataT, numElements> operator&=(const vec<dataT, numElements> &rhs)</code>	Bitwise and of each element of the current vector in-place by the respective element of rhs.
<code>vec<dataT, numElements> operator&=(const dataT &rhs)</code>	Bitwise AND in-place each element of the current vector by rhs.
<code>vec<dataT, numElements> operator%=(const vec<dataT, numElements> &rhs)</code>	Remainder of each element of the current vector in-place by the respective element of rhs.
<code>vec<dataT, numElements> operator%=(const dataT &rhs)</code>	Remainder in-place each element of the current vector by rhs.
<code>vec<dataT, numElements> &operator=(const vec<dataT, numElements> &rhs)</code>	Update each element of the current vector with the respective element of rhs and return a reference to the current vector.
<code>vec<dataT, numElements> &operator=(const dataT &rhs)</code>	Update each element of the current vector with rhs and return a reference to the current vector.
<code>bool operator==(const vec<dataT, numElements> &rhs) const</code>	Return true if all elements of rhs compare equal to the respective element of the current vector.
<code>bool operator!=(const vec<dataT, numElements> &rhs) const</code>	Return true if any one element of rhs does not compare equal to the respective element of the current vector.
End of table	

Table 3.76: Methods for the vec class

3.8 Synchronization and atomics

The SYCL specification offers the same set of synchronization operations that are available to OpenCL C programs, for compatibility and portability across OpenCL devices. 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: OpenCL 1.2 devices only support the equivalent of relaxed C++ atomics and SYCL uses the C++11 library syntax to make this available. This is provided for forward compatibility with future SYCL versions.
- Barriers: Barrier primitives are made available to synchronize sets of work-items within individual work-groups. They are exposed through the `nd_item` class that abstracts the current point in the overall iteration space.
- 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 barrier operations.

Barriers may provide ordering semantics over the local address space, global address space or both. All memory operations initiated before the barrier in the specified address space(s) will be completed before any memory operation after the barrier. Address spaces are described using the `fence_space` enum class:

```
namespace cl {
namespace sycl {
namespace access {
enum class fence_space : char {
    local_space,
    global_space,
    global_and_local
}; // enum class address_space
} // namespaces access
} // namespace sycl
} // namespace cl
```

The SYCL specification provides atomic operations based on the C++11 library syntax. The only available ordering, due to constraints of the OpenCL 1.2 memory model, is `memory_order_relaxed`. No default order is supported because a default order would imply sequential consistency. The SYCL atomic library may map directly to the underlying C++11 library in host code, and must interact safely with the host C++11 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 1.2 device compilers should give a compilation error on use of the `std::atomic` classes and functions in device code. Only `atomic<int>`, `atomic<unsigned int>` and `atomic<float>` types are available in SYCL 1.2. Only the exchange operation is available for float atomics.

No construction of atomic objects is possible in SYCL 1.2. All atomic objects must be obtained by-reference from an accessor (see Section 3.4.6.3).

The atomic types are defined as follows, and methods are listed in Table 3.77:

```
namespace cl {
```

```

namespace sycl {
template <typename T>
class atomic<T> {
public:
    // Constructors
    atomic() = delete;

    // Methods
    // Only memory_order_relaxed is supported in SYCL 1.2
    void store(T operand, std::memory_order = std::memory_order_relaxed);
    void store(T operand, std::memory_order = std::memory_order_relaxed) volatile;
    T load(memory_order = std::memory_order_relaxed) const;
    T load(memory_order = std::memory_order_relaxed) const volatile;

    T exchange(T operand, std::memory_order = std::memory_order_relaxed);
    T exchange(T operand, std::memory_order = std::memory_order_relaxed) volatile;

    T compare_exchange_strong(T* expected, T desired, std::memory_order success,
                               std::memory_order fail);
    T compare_exchange_strong(T* expected, T desired, std::memory_order success,
                               std::memory_order fail) volatile;

    T fetch_add(T operand, std::memory_order = std::memory_order_relaxed);
    T fetch_add(T operand,
                std::memory_order = std::memory_order_relaxed) volatile;

    T fetch_sub(T operand, std::memory_order = std::memory_order_relaxed);
    T fetch_sub(T operand,
                std::memory_order = std::memory_order_relaxed) volatile;

    T fetch_and(T operand, std::memory_order = std::memory_order_relaxed);
    T fetch_and(T operand,
                std::memory_order = std::memory_order_relaxed) volatile;

    T fetch_or(T operand, std::memory_order = std::memory_order_relaxed);
    T fetch_or(T operand, std::memory_order = std::memory_order_relaxed) volatile;

    T fetch_xor(T operand, std::memory_order = std::memory_order_relaxed);
    T fetch_xor(T operand,
                std::memory_order = std::memory_order_relaxed) volatile;

    // Additional functionality provided beyond that of C++11
    T fetch_min(T operand, std::memory_order = std::memory_order_relaxed);
    T fetch_min(T operand,
                std::memory_order = std::memory_order_relaxed) volatile;

    T fetch_max(T operand, std::memory_order = std::memory_order_relaxed);
    T fetch_max(T operand,
                std::memory_order = std::memory_order_relaxed) volatile;
};

typedef atomic<int> atomic_int;
typedef atomic<unsigned int> atomic_uint;
typedef atomic<float> atomic_float;
} // namespace sycl

```

```

} // namespace cl

```

As well as the methods, a matching set of operations on atomic types is provided by the SYCL library. As in the previous case, the only available memory order in SYCL 1.2 is `memory_order_relaxed`. The global functions are as follows and described in Table 3.78.

```

namespace cl {
namespace sycl {
template <class T>
T atomic_load_explicit(atomic<T>* object,
                      std::memory_order = std::memory_order_relaxed);
template <class T>
T atomic_load_explicit(volatile atomic<T>* object,
                      std::memory_order = std::memory_order_relaxed);

template <class T>
void atomic_store_explicit(atomic<T>* object, T,
                          std::memory_order = std::memory_order_relaxed);
template <class T>
void atomic_store_explicit(volatile atomic<T>* object, T operand,
                          std::memory_order = std::memory_order_relaxed);

template <class T>
T atomic_exchange_explicit(atomic<T>* object, T operand,
                          std::memory_order = std::memory_order_relaxed);
template <class T>
T atomic_exchange_explicit(volatile atomic<T>* object, T,
                          std::memory_order = std::memory_order_relaxed);

template <class T>
bool atomic_compare_exchange_strong_explicit(atomic<T>* object, T* expected,
                                             T desired,
                                             std::memory_order success,
                                             std::memory_order fail);
template <class T>
bool atomic_compare_exchange_strong_explicit(volatile atomic<T>* object, T* expected, T desired,
                                             std::memory_order success,
                                             std::memory_order fail);

template <class T>
T atomic_fetch_add_explicit(atomic<T>* object, T operand,
                          std::memory_order = std::memory_order_relaxed);
template <class T>
T atomic_fetch_add_explicit(volatile atomic<T>* object, T operand,
                          std::memory_order = std::memory_order_relaxed);

template <class T>
T atomic_fetch_sub_explicit(atomic<T>* object, T operand,
                          std::memory_order = std::memory_order_relaxed);
template <class T>
T atomic_fetch_sub_explicit(volatile atomic<T>* object, T operand,
                          std::memory_order = std::memory_order_relaxed);

template <class T>

```

```

T atomic_fetch_and_explicit(atomic<T>* object, T operand,
                           std::memory_order = std::memory_order_relaxed);
template <class T>
T atomic_fetch_and_explicit(volatile atomic<T>* object, T operand,
                           std::memory_order = std::memory_order_relaxed);

template <class T>
T atomic_fetch_or_explicit(atomic<T>* object, T operand, std::memory_order);
template <class T>
T atomic_fetch_or_explicit(volatile atomic<T>* object, T operand,
                           std::memory_order = std::memory_order_relaxed);

template <class T>
T atomic_fetch_xor_explicit(atomic<T>* object, T operand,
                           std::memory_order = std::memory_order_relaxed);
template <class T>
T atomic_fetch_xor_explicit(volatile atomic<T>* object, T operand,
                           std::memory_order = std::memory_order_relaxed);

// Additional functionality beyond that provided by C++11
template <class T>
T atomic_fetch_min_explicit(atomic<T>* object, T operand,
                           std::memory_order = std::memory_order_relaxed);
template <class T>
T atomic_fetch_min_explicit(volatile atomic<T>* object, T operand,
                           std::memory_order = std::memory_order_relaxed);

template <class T>
T atomic_fetch_max_explicit(atomic<T>* object, T operand, std::memory_order);
template <class T>
T atomic_fetch_max_explicit(volatile atomic<T>* object, T operand,
                           std::memory_order);
} // namespace sycl
} // namespace cl

```

The atomic operations and methods behave as described in the C++11 specification, barring the restrictions discussed above. Note that care must be taken when using `compare_exchange_strong` to perform many of the operations that would be expected of it in standard CPU code 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.

Methods	Description
<code>void store(T operand, std::memory_order = std::memory_order_relaxed);</code>	Atomically store operand in <code>*this</code> . T may be int, unsigned int or float. order must be <code>memory_order_relaxed</code> .
<code>T load(memory_order = std::memory_order_relaxed) const;</code>	Atomically load the current value of <code>*this</code> and return the value before the call. T may be int, unsigned int or float. order must be <code>memory_order_relaxed</code> .

Continued on next page

Table 3.77: Methods available on an object of type `atomic<T>`.

Methods	Description
<code>T exchange(T operand, std::memory_order = std::memory_order_relaxed);</code>	Atomically replace <code>*this</code> with operand. Return the original value of object. T may be int, unsigned int or float. order must be <code>memory_order_relaxed</code> .
<code>T compare_exchange_strong(T& expected, T desired, std::memory_order success, std::memory_order fail);</code>	Atomically compare the value of <code>*this</code> against expected. If equal replace <code>*this</code> with desired otherwise store the original value of <code>*this</code> in <code>*expected</code> . Returns true if the comparison succeeded. Both memory orders must be <code>memory_order_relaxed</code> . T must be int or unsigned int.
<code>T fetch_add(T operand, std::memory_order = std::memory_order_relaxed);</code>	Atomically add operand to <code>*this</code> . Store the result in <code>*this</code> and return the value before the call. order must be <code>memory_order_relaxed</code> . T must be int or unsigned int.
<code>T fetch_sub(T operand, std::memory_order order);</code>	Atomically subtract operand from <code>*this</code> . Store the result in <code>*this</code> and return the value before the call. order must be <code>memory_order_relaxed</code> . T must be int or unsigned int.
<code>T fetch_and(T operand, std::memory_order order);</code>	Atomically perform a bitwise and of operand and <code>*this</code> . Store the result in <code>*this</code> and return the value before the call. order must be <code>memory_order_relaxed</code> . T must be int or unsigned int.
<code>T fetch_or(T operand, std::memory_order order);</code>	Atomically perform a bitwise or of operand and <code>*this</code> . Store the result in <code>*this</code> and return the value before the call. order must be <code>memory_order_relaxed</code> . T must be int or unsigned int.
<code>T fetch_xor(T operand, std::memory_order order);</code>	Atomically perform a bitwise exclusive-or of operand and <code>*this</code> . Store the result in <code>*this</code> and return the value before the call. order must be <code>memory_order_relaxed</code> . T must be int or unsigned int.
<code>T fetch_min(T operand, std::memory_order order);</code>	Atomically compute the minimum of operand and <code>*this</code> . Store the result in <code>*this</code> and return the value before the call. order must be <code>memory_order_relaxed</code> . T must be int or unsigned int.
<code>T fetch_max(T operand, std::memory_order order);</code>	Atomically compute the maximum of operand and <code>*this</code> . Store the result in <code>*this</code> and return the value before the call. order must be <code>memory_order_relaxed</code> . T must be int or unsigned int.
End of table	

Table 3.77: Methods available on an object of type `atomic<T>`.

Functions	Description
<code>template<class T> T atomic_load_explicit(atomic<T>* object, std::memory_order order);</code>	Atomically load the current value of object and return that value. T may be int, unsigned int or float. order must be memory_order_relaxed.
<code>template<class T> void atomic_store_explicit(atomic<T>* object, T operand, std::memory_order order);</code>	Atomically store operand in object. T may be int, unsigned int or float. order must be memory_order_relaxed.
<code>template<class T> T atomic_exchange_explicit(atomic<T>* object, T operand, std::memory_order order);</code>	Atomically replace object with operand. Return the original value of object. T may be int, unsigned int or float. order must be memory_order_relaxed.
<code>template<class T> bool atomic_compare_exchange_strong_explicit(atomic<T>* object, T* expected, T desired, std::memory_order success, std::memory_order fail);</code>	Atomically compare the value of object against expected. If equal replace object with desired. Otherwise store the original value of object in expected. Returns true if the comparison succeeded. Both memory orders must be memory_order_relaxed. T must be int or unsigned int.
<code>template<class T> T atomic_fetch_add_explicit(atomic<T>* object, T operand, std::memory_order order);</code>	Atomically add operand to object. Store the result in object. order must be memory_order_relaxed. T must be int or unsigned int.
<code>template<class T> T atomic_fetch_sub_explicit(atomic<T>* object, T operand, std::memory_order order);</code>	Atomically subtract operand from object. Store the result in object. order must be memory_order_relaxed. T must be int or unsigned int.
<code>template<class T> T atomic_fetch_and_explicit(atomic<T>* operand, T object, std::memory_order order);</code>	Atomically perform a bitwise and of operand and object. Store the result in object. order must be memory_order_relaxed. T must be int or unsigned int.
<code>template<class T> T atomic_fetch_or_explicit(atomic<T>* object, T operand, std::memory_order order);</code>	Atomically perform a bitwise or of operand and object. Store the result in object. order must be memory_order_relaxed. T must be int or unsigned int.
<code>template<class T> T atomic_fetch_xor_explicit(atomic<T>* object, T operand, std::memory_order order);</code>	Atomically perform a bitwise exclusive-or of operand and object. Store the result in object. order must be memory_order_relaxed. T must be int or unsigned int.
<code>template<class T> T atomic_fetch_min_explicit(atomic<T>* object, T operand, std::memory_order order);</code>	Atomically compute the minimum of operand and object. Store the result in object. order must be memory_order_relaxed. T must be int or unsigned int.

Continued on next page

Table 3.78: Global functions available on atomic types.

Functions	Description
<pre>template<class T> T atomic_fetch_max_explicit(atomic<T>* object, T operand, std::memory_order order);</pre>	Atomically compute the maximum of operand and object. Store the result in object. order must be memory_order_relaxed. T must be int or unsigned int.
End of table	

Table 3.78: Global functions available on atomic types.

3.9 Stream class

SYCL stream class is available in SYCL instead of the `printf()` function in order to support C++ classes and structs and some commonly used stream output manipulators used in C++.

The `cl::sycl::stream` object can be used in order to output a sequence of characters on standard output. The stream object can be default constructed inside a command group scope and it can also have implementation defined constructors for output size and width.

Stream Operators	Description
<code>cl::sycl::stream=</code>	operator for assigning a stream to another.
<code>cl::sycl::operator<<</code>	operator for outputting a sequence of character or structs on standard output. All SYCL vector classes and id structs can be passed in this operator.
<code>cl::sycl::endl</code>	method that adds an end of line character to the output
<code>cl::sycl::hex</code>	Display the output in hexadecimal base
<code>cl::sycl::oct</code>	Display the output in octal base
<code>cl::sycl::setw</code>	Sets the field width of the output
<code>cl::sycl::precision</code>	Sets the decimal precision for the output
<code>cl::sycl::scientific</code>	Set the notation of the floating-point values as the C++ scientific notation
<code>cl::sycl::fixed</code>	Set the notation of the floating-point values as the C++ fixed notation
<code>cl::sycl::hexfloat</code>	Displays the floating point values in hexadecimal format
<code>cl::sycl::defaultfloat</code>	Displays the floating point values in the default notation.

The usage of the `cl::sycl::stream` object is not recommend for performance critical applications, as optimization levels for streaming operations are implementation defined, and the corresponding implementation documentation should be consulted.

3.10 SYCL built-in functions for SYCL host and device

SYCL kernels may execute on any SYCL device- OpenCL device or SYCL host, which requires that the functions used in the kernels to be compiled and linked for both device and host. In the SYCL system the OpenCL built-ins are available for the SYCL host and device within the `cl::sycl` namespace, although, their semantics may be different. This section follows the OpenCL 1.2 specification document [1, ch. 6.12] and describes the behavior of these functions for SYCL host and device.

The SYCL built-in functions are available throughout the SYCL application, and depending 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.

3.10.1 Description of the built-in types available for SYCL host and device

All the OpenCL built-in types are available in the namespace `cl::sycl`. For the purposes of this document we use type names for describing sets of SYCL valid types. The type names themselves are not valid SYCL types, but they represent a set of valid types, as defined in tables 3.80.

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 3.5.1.4, 3.5.1.5 and 3.5.1.6.

Generic type name	Description
<code>floatn</code>	<code>cl::sycl::float2</code> , <code>cl::sycl::float3</code> , <code>cl::sycl::float4</code> , <code>cl::sycl::float8</code> , <code>cl::sycl::float16</code>
<code>genfloatf</code>	<code>float</code> , <code>floatn</code>
<code>doublen</code>	<code>cl::sycl::double2</code> , <code>cl::sycl::double3</code> , <code>cl::sycl::double4</code> , <code>cl::sycl::double8</code> , <code>cl::sycl::double16</code>
<code>genfloatd</code>	<code>double</code> , <code>doublen</code>
<code>genfloat</code>	<code>float</code> , <code>floatn</code> <code>double</code> , <code>doublen</code>
<code>sgenfloat</code>	<code>float</code> , <code>double</code> It is a scalar type that matches the corresponding vector type <code>floatn</code> or <code>doublen</code> .
<code>charn</code>	<code>cl::sycl::char2</code> , <code>cl::sycl::char3</code> , <code>cl::sycl::char4</code> , <code>cl::sycl::char8</code> , <code>cl::sycl::char16</code>
<code>ucharn</code>	<code>cl::sycl::uchar2</code> , <code>cl::sycl::uchar3</code> , <code>cl::sycl::uchar4</code> , <code>cl::sycl::uchar8</code> , <code>cl::sycl::uchar16</code>
<code>genchar</code>	<code>char</code> , <code>charn</code>
<code>ugenchar</code>	<code>unsigned char</code> , <code>ugenchar</code>

Continued on next page

Table 3.80: Generic type name description, which serves as a description for all valid types of parameters to kernel functions. [1]

Generic type name	Description
shortn	cl::sycl::short2, cl::sycl::short3, cl::sycl::short4, cl::sycl::short8, cl::sycl::short16
genshort	short, shortn
ushortn	cl::sycl::ushort2, cl::sycl::ushort3, cl::sycl::ushort4, cl::sycl::ushort8, cl::sycl::ushort16
ugenshort	unsigned short, ugenshort
uintn	cl::sycl::uint2, cl::sycl::uint3, cl::sycl::uint4, cl::sycl::uint8, cl::sycl::uint16
ugenint	int, uintn
intn	cl::sycl::int2, cl::sycl::int3, cl::sycl::int4, cl::sycl::int8, cl::sycl::int16
genint	int, intn
ulonglongn	cl::sycl::ulonglong2, cl::sycl::ulonglong3, cl::sycl::ulonglong4, cl::sycl::ulonglong8, cl::sycl::ulonglong16
ugenlonglong	unsigned long long int, ulonglongn
longlongn	cl::sycl::longlong2, cl::sycl::longlong3, cl::sycl::longlong4, cl::sycl::longlong8, cl::sycl::longlong16
genlonglong	long long int, longlongn
geninteger	genchar, ugenchar, genshort, ugenshort, genint, ugenint, genlonglong, ugenlonglong
sgeninteger	char, short, int, long long int, unsigned char, unsigned short, unsigned int, unsigned long long int
ugeninteger	uchar, ucharn, ushort, ushortn, uint, uintn, ulonglong, ulonglongn
gentype	char, charn, uchar, ucharn, short, shortn, ushort, ushortn, int, intn, uint, uintn, longlong, longlongn, ulonglong, ulonglongn, float, floatn, double, doublen.
End of table	

Table 3.80: Generic type name description, which serves as a description for all valid types of parameters to kernel functions. [1]

3.10.2 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 `nd_item` and `group` classes see section 3.5.1.5 and 3.5.1.6.

3.10.3 Math functions

In SYCL the OpenCL math functions are available in the namespace `cl::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* counterparts, for dimensions 2, 3, 4, 8 and 16. On the host the vector types are going to be using the *vec* class and on an OpenCL device are going to be using the corresponding OpenCL vector types.

The built-in functions available for SYCL host and device with the same precision requirements for both host and device, are described in table 3.81.

Math Function	Description
<code>genfloat acos (genfloat x)</code>	Inverse cosine function.
<code>genfloat acosh (genfloat x)</code>	Inverse hyperbolic cosine.
<code>genfloat acospi (genfloat x)</code>	Compute $\text{acos}x/\pi$
<code>genfloat asin (genfloat x)</code>	Inverse sine function.
<code>genfloat asinh (genfloat x)</code>	Inverse hyperbolic sine.
<code>genfloat asinpi (genfloat x)</code>	Compute $\text{asin}x/\pi$
<code>genfloat atan (genfloat y_over_x)</code>	Inverse tangent function.
<code>genfloat atan2 (genfloat y, genfloat x)</code>	Compute $\text{atan}(y/x)$.
<code>genfloat atanh (genfloat x)</code>	Hyperbolic inverse tangent.
<code>genfloat atanpi (genfloat x)</code>	Compute $\text{atan}(x)/\pi$.
<code>genfloat atan2pi (genfloat y, genfloat x)</code>	Compute $\text{atan2}(y, x)/\pi$.
<code>genfloat cbrt (genfloat x)</code>	Compute cube-root.
<code>genfloat ceil (genfloat x)</code>	Round to integral value using the round to positive infinity rounding mode.
<code>genfloat copysign (genfloat x, genfloat y)</code>	Returns x with its sign changed to match the sign of y .
<code>genfloat cos (genfloat x)</code>	Compute cosine.
<code>genfloat cosh (genfloat x)</code>	Compute hyperbolic cosine.
<code>genfloat cospi (genfloat x)</code>	Compute $\cos(\pi x)$.
<code>genfloat erfc (genfloat x)</code>	Complementary error function.
<code>genfloat erf (genfloat x)</code>	Error function encountered in integrating the normal distribution.
<code>genfloat exp (genfloat x)</code>	Compute the base-e exponential of x .
<code>genfloat exp2 (genfloat x)</code>	Exponential base 2 function.
<code>genfloat exp10 (genfloat x)</code>	Exponential base 10 function.
<code>genfloat expm1 (genfloat x)</code>	Compute $\exp(x) - 1.0$.

Continued on next page

Table 3.81: Math functions which work on SYCL Host and device. They correspond to Table 6.7 of the OpenCL 1.2 specification [1]

Math Function	Description
<code>genfloat fabs (genfloat x)</code>	Compute absolute value of a floating-point number.
<code>genfloat fdim (genfloat x, genfloat y)</code>	$x - y$ if $x > y$, +0 if x is less than or equal to y .
<code>genfloat floor (genfloat x)</code>	Round to integral value using the round to negative infinity rounding mode.
<code>genfloat fma (genfloat a, genfloat b, genfloat c)</code>	Returns the correctly rounded floating-point representation of the sum of c with the infinitely precise product of a and b . Rounding of intermediate products shall not occur. Edge case behavior is per the IEEE 754-2008 standard.
<code>genfloat fmax (genfloat x, genfloat y)</code> <code>genfloat fmax (genfloat x, sgenfloat y)</code>	Returns y if $x < y$, otherwise it returns x . If one argument is a NaN, <code>fmax()</code> returns the other argument. If both arguments are NaNs, <code>fmax()</code> returns a NaN.
<code>genfloat fmin (genfloat x, genfloat y)</code> <code>genfloat fmin (genfloat x, sgenfloat y)</code>	Returns y if $y < x$, otherwise it returns x . If one argument is a NaN, <code>fmin()</code> returns the other argument. If both arguments are NaNs, <code>fmin()</code> returns a NaN.
<code>genfloat fmod (genfloat x, genfloat y)</code>	Modulus. Returns $xy * trunc(x/y)$.
<code>floatn fract (floatn x, intn *iptr)</code> <code>float fract (float x, int * iptr)</code>	Returns <code>fmin(x - floor(x), 0x1.ffffep-1f)</code> . <code>floor(x)</code> is returned in <code>iptr</code> .
<code>doublen frexp (doublen x, intn *exp)</code> <code>double frexp (double x, int * exp)</code>	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 * 2^{exp} .
<code>genfloat hypot (genfloat x, genfloat y)</code>	Compute the value of the square root of $x^2 + y^2$ without undue overflow or underflow.
<code>int logb (float x)</code> <code>intn ilogb (genfloat x)</code> <code>int logb (double x)</code> <code>intn logb (doublen x)</code>	Return the exponent as an integer value.
<code>genfloat ldexp (genfloat x, genint k)</code> <code>floatn ldexp (floatn x, int k)</code> <code>doublen ldexp (doublen x, int k)</code>	Multiply x by 2 to the power k .
<code>genfloat lgamma (genfloat x)</code>	Log gamma function. Returns the natural logarithm of the absolute value of the gamma function. The sign of the gamma function is returned in the <code>signp</code> argument of <code>lgamma_r</code> .
<code>genfloat lgamma_r (genfloat x, genint *signp)</code>	Log gamma function. Returns the natural logarithm of the absolute value of the gamma function. The sign of the gamma function is returned in the <code>signp</code> argument of <code>lgamma_r</code> .

Continued on next page

Table 3.81: Math functions which work on SYCL Host and device. They correspond to Table 6.7 of the OpenCL 1.2 specification [1]

Math Function	Description
<code>genfloat log (genfloat)</code>	Compute natural logarithm.
<code>genfloat log2 (genfloat)</code>	Compute a base 2 logarithm.
<code>genfloat log10 (genfloat)</code>	Compute a base 10 logarithm.
<code>genfloat log1p (genfloat x)</code>	Compute $\log_e(1.0 + x)$.
<code>genfloat logb (genfloat x)</code>	Compute the exponent of x , which is the integral part of $\log_r(x)$.
<code>genfloat mad (genfloat a, genfloat b, genfloat c)</code>	<code>mad</code> 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. <code>mad</code> is intended to be used where speed is preferred over accuracy.
<code>genfloat maxmag (genfloat x, genfloat y)</code>	Returns x if $ x > y $, y if $ y > x $, otherwise $\text{fmax}(x, y)$.
<code>genfloat minmag (genfloat x, genfloat y)</code>	Returns x if $ x < y $, y if $ y < x $, otherwise $\text{fmin}(x, y)$.
<code>genfloat modf (genfloat x, genfloat *iptr)</code>	Decompose a floating-point number. The <code>modf</code> function breaks the argument x into integral and fractional parts, each of which has the same sign as the argument. It stores the integral part in the object pointed to by <code>iptr</code> .
<code>floatn nan (uintn nancode)</code> <code>float nan (unsigned int nancode)</code> <code>doublen nan (ulonglongn nancode)</code> <code>double nan (unsigned long long int nancode)</code>	Returns a quiet NaN. The nancode may be placed in the significand of the resulting NaN.
<code>genfloat nextafter (genfloat x, genfloat y)</code>	Computes the next representable single-precision floating-point value following x in the direction of y . Thus, if y is less than x , <code>nextafter()</code> returns the largest representable floating-point number less than x .
<code>genfloat pow (genfloat x, genfloat y)</code>	Compute x to the power y .
<code>genfloat pown (genfloat x, genint y)</code>	Compute x to the power y , where y is an integer.
<code>genfloat powr (genfloat x, genfloat y)</code>	Compute x to the power y , where $x \geq 0$.
<code>genfloat remainder (genfloat x, genfloat y)</code>	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 .

Continued on next page

Table 3.81: Math functions which work on SYCL Host and device. They correspond to Table 6.7 of the OpenCL 1.2 specification [1]

Math Function	Description
<code>genfloat remquo (genfloat x, genfloat y, genint *quo)</code>	The <code>remquo</code> function computes the value <code>r</code> such that $r = x - k*y$, where <code>k</code> is the integer nearest the exact value of x/y . If there are two integers closest to x/y , <code>k</code> shall be the even one. If <code>r</code> is zero, it is given the same sign as <code>x</code> . This is the same value that is returned by the remainder function. <code>remquo</code> also calculates the lower seven bits of the integral quotient x/y , and gives that value the same sign as x/y . It stores this signed value in the object pointed to by <code>quo</code> .
<code>genfloat rint (genfloat)</code>	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.
<code>genfloat rootn (genfloat x, genint y)</code>	Compute x to the power $1/y$.
<code>genfloat round (genfloat x)</code>	Return the integral value nearest to <code>x</code> rounding halfway cases away from zero, regardless of the current rounding direction.
<code>genfloat rsqrt (genfloat)</code>	Compute inverse square root.
<code>genfloat sin (genfloat)</code>	Compute sine.
<code>genfloat sincos (genfloat x, genfloat *cosval)</code>	Compute sine and cosine of <code>x</code> . The computed sine is the return value and computed cosine is returned in <code>cosval</code> .
<code>genfloat sinh (genfloat x)</code>	Compute hyperbolic sine.
<code>genfloat sinpi (genfloat x)</code>	Compute $\sin(\pi x)$.
<code>genfloat sqrt (genfloat x)</code>	Compute square root.
<code>genfloat tan (genfloat x)</code>	Compute tangent.
<code>genfloat tanh (genfloat x)</code>	Compute hyperbolic tangent.
<code>genfloat tanpi (genfloat x)</code>	Compute $\tan(\pi x)$.
<code>genfloat tgamma (genfloat x)</code>	Compute the gamma function.
<code>genfloat trunc (genfloat x)</code>	Round to integral value using the round to zero rounding mode.
End of table	

Table 3.81: Math functions which work on SYCL Host and device. They correspond to Table 6.7 of the OpenCL 1.2 specification [1]

In SYCL the implementation defined precision math functions are defined in the namespace `cl::sycl::native`. The functions that are available within this namespace are specified in tables 3.82.

Native Math Function	Description
<code>genfloat cos (genfloat x)</code>	Compute cosine over an implementation-defined range. The maximum error is implementation-defined.
Continued on next page	

Table 3.82: Native functions which work on SYCL Host and device, are available in the `cl::sycl::native` namespace. They correspond to Table 6.9 of the OpenCL 1.2 specification [1]

Native Math Function	Description
<code>genfloat divide (genfloat x, genfloat y)</code>	Compute x / y over an implementation-defined range. The maximum error is implementation-defined.
<code>genfloat exp (genfloat x)</code>	Compute the base- e exponential of x over an implementation-defined range. The maximum error is implementation-defined.
<code>genfloat exp2 (genfloat x)</code>	Compute the base- 2 exponential of x over an implementation-defined range. The maximum error is implementation-defined.
<code>genfloat exp10 (genfloat x)</code>	Compute the base- 10 exponential of x over an implementation-defined range. The maximum error is implementation-defined.
<code>genfloat log (genfloat x)</code>	Compute natural logarithm over an implementation defined range. The maximum error is implementation-defined.
<code>genfloat log2 (genfloat x)</code>	Compute a base 2 logarithm over an implementation-defined range. The maximum error is implementation-defined.
<code>genfloat log10 (genfloat x)</code>	Compute a base 10 logarithm over an implementation-defined range. The maximum error is implementation-defined.
<code>genfloat powr (genfloat x, genfloat y)</code>	Compute x to the power y , where $x \geq 0$. The range of x and y are implementation-defined. The maximum error is implementation-defined.
<code>genfloat recip (genfloat x)</code>	Compute reciprocal over an implementation-defined range. The maximum error is implementation-defined.
<code>genfloat rsqrt (genfloat x)</code>	Compute inverse square root over an implementation-defined range. The maximum error is implementation-defined.
<code>genfloat sin (genfloat x)</code>	Compute sine over an implementation-defined range. The maximum error is implementation-defined.
<code>genfloat sqrt (genfloat x)</code>	Compute square root over an implementation-defined range. The maximum error is implementation-defined.
<code>genfloat tan (genfloat x)</code>	Compute tangent over an implementation-defined range. The maximum error is implementation-defined.
End of table	

Table 3.82: Native functions which work on SYCL Host and device, are available in the `cl::sycl::native` namespace. They correspond to Table 6.9 of the OpenCL 1.2 specification [1]

3.10.4 Integer functions

In SYCL the OpenCL integer math functions are available in the namespace `cl::sycl` on host and device as defined in the OpenCL 1.2 specification document [1, par. 6.12.3]. 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` counterparts, for dimensions 2, 3, 4, 8 and 16. On the host the vector types are going to be using the `vec` class and on an OpenCL device are going to be using the corresponding OpenCL vector types. The supported integer math functions are described in tables 3.83.

Integer Function	Description
<code>ugeninteger abs (geninteger x)</code>	Returns $ x $.
<code>ugeninteger abs_diff (geninteger x, geninteger y)</code>	Returns $ x - y $ without modulo overflow.
<code>geninteger add_sat (geninteger x, geninteger y)</code>	Returns $x + y$ and saturates the result.
<code>geninteger hadd (geninteger x, geninteger y)</code>	Returns $(x + y) \gg 1$. The intermediate sum does not modulo overflow.
<code>geninteger rhadd (geninteger x, geninteger y)</code>	Returns $(x + y + 1) \gg 1$. The intermediate sum does not modulo overflow.
<code>geninteger clamp (geninteger x, sgeninteger minval, sgeninteger maxval)</code>	Returns $\min(\max(x, \text{minval}), \text{maxval})$. Results are undefined if $\text{minval} > \text{maxval}$.
<code>geninteger clz (geninteger x)</code>	Returns the number of leading 0-bits in x , starting at the most significant bit position.
<code>geninteger clamp (geninteger x, geninteger minval, geninteger maxval)</code>	Returns $\min(\max(x, \text{minval}), \text{maxval})$. Results are undefined if $\text{minval} > \text{maxval}$.
<code>geninteger mad_hi (geninteger a, geninteger b, geninteger c)</code>	Returns $\text{mul_hi}(a, b) + c$.
<code>geninteger mad_sat (geninteger a, geninteger b, geninteger c)</code>	Returns $a * b + c$ and saturates the result.
<code>geninteger max (geninteger x, geninteger y)</code> <code>geninteger max (geninteger x, sgeninteger y)</code>	Returns y if $x < y$, otherwise it returns x .
<code>geninteger min (geninteger x, geninteger y)</code> <code>geninteger min (geninteger x, sgeninteger y)</code>	Returns y if $y < x$, otherwise it returns x .
<code>geninteger mul_hi (geninteger x, geninteger y)</code>	Computes $x * y$ and returns the high half of the product of x and y .
<code>geninteger rotate (geninteger v, geninteger i)</code>	For each element in v , the bits are shifted left by the number of bits given by the corresponding element in i (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.
<code>geninteger sub_sat (geninteger x, geninteger y)</code>	Returns $x - y$ and saturates the result.
<code>shortn upsample (charn hi, uchar n lo)</code>	$result[i] = ((short)hi[i] \ll 8) lo[i]$
<code>ushortn upsample (uchar n hi, uchar n lo)</code>	$result[i] = ((ushort)hi[i] \ll 8) lo[i]$
<code>intn upsample (shortn hi, ushortn lo)</code>	$result[i] = ((int)hi[i] \ll 16) lo[i]$
<code>uintn upsample (ushortn hi, ushortn lo)</code>	$result[i] = ((uint)hi[i] \ll 16) lo[i]$
<code>longlongn upsample (intn hi, uintn lo)</code>	$result[i] = ((long)hi[i] \ll 32) lo[i]$
<code>ulonglongn upsample (uintn hi, uintn lo)</code>	$result[i] = ((ulong)hi[i] \ll 32) lo[i]$
<code>geninteger popcount (geninteger x)</code>	Returns the number of non-zero bits in x .

Continued on next page

Table 3.83: Integer functions which work on SYCL Host and device, are available in the `cl::sycl` namespace. They correspond to Table 6.10 of the OpenCL 1.2 specification [1]

Integer Function	Description
<pre>intn mad24 (intn x, intn y, intn z) uintn mad24 (uintn x, uintn y, uintn z)</pre>	Multiply two 24-bit integer values x and y 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>intn mul24 (intn x, intn y) uintn mul24 (uintn x, uintn y)</pre>	Multiply two 24-bit integer values x and y. x and y are 32-bit integers but only the low 24-bits are used to perform the multiplication. mul24 should only be used when values in x and y are in the range [- 223, 223-1] if x and y are signed integers and in the range [0, 224-1] if x and y are unsigned integers. If x and y are not in this range, the multiplication result is implementation-defined.
End of table	

Table 3.83: Integer functions which work on SYCL Host and device, are available in the `cl::sycl` namespace. They correspond to Table 6.10 of the OpenCL 1.2 specification [1]

3.10.5 Common functions

In SYCL the OpenCL *common* functions are available in the namespace `cl::sycl` on host and device as defined in the OpenCL 1.2 specification document [1, par. 6.12.4]. Description is in table 3.84. The built-in functions can take as input float or optionally double and their *vec* counterparts, for dimensions 2, 3, 4, 8 and 16. On the host the vector types are going to be using the *vec* class and on an OpenCL device are going to be using the corresponding OpenCL vector types.

Common Function	Description
<pre>genfloat clamp (genfloat x, genfloat minval, genfloat maxval) floatn clamp (floatn x, float minval, float maxval)) doublen clamp (doublen x, double minval, double maxval)</pre>	Returns $f_{\min}(f_{\max}(x, \text{minval}), \text{maxval})$. Results are undefined if $\text{minval} > \text{maxval}$.
<pre>genfloat degrees (genfloat radians)</pre>	Converts radians to degrees, i.e. $(180/\pi) * \text{radians}$.
<pre>genfloat max (genfloat x, genfloat y) genfloatf max (genfloatf x, float y) genfloatd max (genfloatd x, double y)</pre>	Returns y if $x < y$, otherwise it returns x. If x or y are infinite or NaN, the return values are undefined.
<pre>genfloat min (genfloat x, genfloat y) genfloatf min (genfloatf x, float y) genfloatd min (genfloatd x, double y)</pre>	Returns y if $y < x$, otherwise it returns x. If x or y are infinite or NaN, the return values are undefined.
Continued on next page	

Table 3.84: Common functions which work on SYCL Host and device, are available in the `cl::sycl` namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1]

Common Function	Description
<code>genfloat mix (genfloat x, genfloat y, genfloat a)</code> <code>genfloatf mix (genfloatf x, genfloatf y, float a)</code> <code>genfloatd mix (genfloatd x, genfloatd y, double a)</code>	Returns the linear blend of x & y implemented as: $x+(y-x)*a$. a must be a value in the range 0.0 ... 1.0. If a is not in the range 0.0 ... 1.0, the return values are undefined.
<code>genfloat radians (genfloat degrees)</code>	Converts degrees to radians, i.e. $(\pi/180) * degrees$.
<code>genfloat step (genfloat edge, genfloat x)</code> <code>genfloatf step (float edge, genfloatf x)</code> <code>genfloatd step (double edge, genfloatd x)</code>	Returns 0.0 if $x < edge$, otherwise it returns 1.0.
<code>genfloat smoothstep (genfloat edge0, genfloat edge1, genfloat x)</code> <code>genfloatf smoothstep (float edge0, float edge1, genfloatf x)</code> <code>genfloatd smoothstep (double edge0, double edge1, genfloatd x)</code>	<p>Returns 0.0 if $x \leq edge0$ and 1.0 if $x \geq edge1$ and performs smooth Hermite interpolation between 0 and 1 when $edge0 < x < edge1$. This is useful in cases where you would want a threshold function with a smooth transition.</p> <p>This is equivalent to:</p> <pre> gentype t; t = clamp ((x <= edge0)/ (edge1 >= edge0), 0, 1); return t * t * (3 - 2 * t); </pre> <p>Results are undefined if $edge0 \geq edge1$ or if x, $edge0$ or $edge1$ is a NaN.</p>
<code>genfloat sign (genfloat x)</code>	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 3.84: Common functions which work on SYCL Host and device, are available in the `cl::sycl` namespace. They correspond to Table 6.12 of the OpenCL 1.2 specification [1]

3.10.6 Geometric Functions

In SYCL the OpenCL *geometric* functions are available in the namespace `cl::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* counterparts, for dimensions 2, 3, 4, 8 and 16. On the host the vector types are going to be using the *vec* class and on an OpenCL device are going to be using the corresponding OpenCL vector types. All of the geometric functions are using round-to-nearest-even rounding mode. Tables 3.85 contain all the definitions of supported geometric functions.

Geometric Function	Description
<code>float4 cross (float4 p0, float4 p1)</code> <code>float3 cross (float3 p0, float3 p1)</code> <code>double4 cross (double4 p0, double4 p1)</code> <code>double3 cross (double3 p0, double3 p1)</code>	Returns the cross product of $p0.xyz$ and $p1.xyz$. The w component of float4 result returned will be 0.0.
Continued on next page	

Table 3.85: Geometric functions which work on SYCL Host and device, are available in the `cl::sycl` namespace. They correspond to Table 6.13 of the OpenCL 1.2 specification [1]₁₆₄

Geometric Function	Description
<code>float dot (floatn p0, floatn p1)</code> <code>double dot (doublen p0, doublen p1)</code>	Compute dot product.
<code>float distance (floatn p0, floatn p1)</code> <code>double distance (doublen p0, doublen p1)</code>	Returns the distance between p0 and p1. This is calculated as $length(p0 - p1)$.
<code>float length (floatn p)</code> <code>double length (doublen p)</code>	Return the length of vector p, i.e., $\sqrt{p.x^2 + p.y^2 + \dots}$
<code>floatn normalize (floatn p)</code> <code>doublen normalize (doublen p)</code>	Returns a vector in the same direction as p but with a length of 1.
<code>float fast_distance (floatn p0, floatn p1)</code>	Returns $fastlength(p0 - p1)$.
<code>float fast_length (floatn p)</code>	Returns the length of vector p computed as: $\text{sqrt}((\text{half})(\text{pow}(p.x,2) + \text{pow}(p.y,2) + \dots))$
<code>floatn fast_normalize (floatn p)</code>	Returns a vector in the same direction as p but with a length of 1. <code>fast_normalize</code> is computed as: $p * \text{rsqrt}((\text{half})(\text{pow}(p.x,2) + \text{pow}(p.y,2) + \dots))$ The result shall be within 8192 ulps error from the infinitely precise result of <code>if</code> (all $(p == 0.0f)$) <code>result = p;</code> <code>else</code> <code>result = p / sqrt (pow(p.x,2)+ pow(p.y,2)+ ...);</code> with the following exceptions: <ol style="list-style-type: none"> 1. If the sum of squares is greater than <code>FLT_MAX</code> then the value of the floating-point values in the result vector are undefined. 2. If the sum of squares is less than <code>FLT_MIN</code> 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 $\text{sqrt}(FLT_MIN)$ may be flushed to zero before proceeding with the calculation.
End of table	

Table 3.85: Geometric functions which work on SYCL Host and device, are available in the `cl::sycl` namespace. They correspond to Table 6.13 of the OpenCL 1.2 specification [1]

3.10.7 Relational functions

In SYCL the OpenCL *relational* functions are available in the namespace `cl::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* counterparts, for dimensions 2,3,4,8, and 16. On the host the vector types are going to be using the *vec* class and on an OpenCL device are going to be using the corresponding OpenCL vector types. The relational operators are available in both host and device, these relational functions are provided in addition to the the operators and will return 0 if the conditional is *false* and 1 otherwise. The available built-in functions are described in tables 3.86

Relational Function	Description
<code>int isequal (float x, float y)</code> <code>intn isequal (floatn x, floatn y)</code> <code>longlong isequal (double x, double y)</code> <code>longlongn isequal (doublen x, doublen y)</code>	Returns the component-wise compare of $x == y$.
<code>int isnotequal (float x, float y)</code> <code>intn isnotequal (floatn x, floatn y)</code> <code>longlong isnotequal (double x, double y)</code> <code>longlongn isnotequal (doublen x, doublen y)</code>	Returns the component-wise compare of $x! = y$.
<code>int isgreater (float x, float y)</code> <code>intn isgreater (floatn x, floatn y)</code> <code>longlong isgreater (double x, double y)</code> <code>longlongn isgreater (doublen x, doublen y)</code>	Returns the component-wise compare of $x > y$.
<code>int isgreaterequal (float x, float y)</code> <code>intn isgreaterequal (floatn x, floatn y)</code> <code>longlong isgreaterequal (double x, double y)</code> <code>longlongn isgreaterequal (doublen x, doublen y)</code>	Returns the component-wise compare of $x >= y$.
<code>int isless (float x, float y)</code> <code>intn isless (floatn x, floatn y)</code> <code>longlong isless (double x, double y)</code> <code>longlongn isless (doublen x, doublen y)</code>	Returns the component-wise compare of $x < y$.
<code>int islessequal (float x, float y)</code> <code>intn islessequal (floatn x, floatn y)</code> <code>longlong islessequal (double x, double y)</code> <code>longlongn islessequal (doublen x, doublen y)</code>	Returns the component-wise compare of $x <= y$.
<code>int islessgreater (float x, float y)</code> <code>intn islessgreater (floatn x, floatn y)</code> <code>longlong islessgreater (double x, double y)</code> <code>longlongn islessgreater (doublen x, doublen y)</code>	Returns the component-wise compare of $(x < y) (x > y)$.
<code>int isfinite (float)</code> <code>intn isfinite (floatn)</code> <code>longlong isfinite (double)</code> <code>longlongn isfinite (doublen)</code>	Test for finite value.
<code>int isinf (float)</code> <code>intn isinf (floatn)</code> <code>longlong isinf (double)</code> <code>longlongn isinf (doublen)</code>	Test for infinity value (positive or negative) .

Continued on next page

Table 3.86: Relational functions which work on SYCL Host and device, are available in the `cl::sycl` namespace. They correspond to Table 6.14 of the OpenCL 1.2 specification [1]

Relational Function	Description
<code>int isnan (float)</code> <code>intn isnan (floatn)</code> <code>longlong isnan (double)</code> <code>longlongn isnan (doublen)</code>	Test for a NaN.
<code>int isnormal (float)</code> <code>intn isnormal (floatn)</code> <code>longlong isnormal (double)</code> <code>longlongn isnormal (doublen)</code>	Test for a normal value.
<code>int isordered (float x, float y)</code> <code>intn isordered (floatn x, floatn y)</code> <code>longlong isordered (double x, double y)</code> <code>longlongn isordered (doublen x, doublen y)</code>	Test if arguments are ordered. <code>isordered()</code> takes arguments <code>x</code> and <code>y</code> , and returns the result <code>isequal(x, x) && isequal(y, y)</code> .
<code>int isunordered (float x, float y)</code> <code>intn isunordered (floatn x, floatn y)</code> <code>longlong isunordered (double x, double y)</code> <code>longlongn isunordered (doublen x, doublen y)</code>	Test if arguments are unordered. <code>isunordered()</code> takes arguments <code>x</code> and <code>y</code> , returning non-zero if <code>x</code> or <code>y</code> is NaN, and zero otherwise.
<code>int signbit (float)</code> <code>intn signbit (floatn)</code> <code>longlong signbit (double)</code> <code>longlongn signbit (doublen)</code>	Test for sign bit. The scalar version of the 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 <code>floatn</code> : -1 (i.e all bits set) if the sign bit in the float is set else returns 0.
<code>int any (geninteger x)</code>	Returns 1 if the most significant bit in any component of <code>x</code> is set; otherwise returns 0.
<code>int all (ugeninteger x)</code>	Returns 1 if the most significant bit in all components of <code>x</code> is set; otherwise returns 0.
<code>gentype bitselect (gentype a, gentype b, gentype c)</code>	Each bit of the result is the corresponding bit of <code>a</code> if the corresponding bit of <code>c</code> is 0. Otherwise it is the corresponding bit of <code>b</code> .
<code>gentype select (gentype a, gentype b, geninteger c)</code>	For each component of a vector type: <code>result[i] = (MSB of c[i] is set)? b[i] : a[i]</code> . For a scalar type: <code>result = c ? b : a</code> . <code>geninteger</code> must have the same number of elements and bits as <code>gentype</code> .
End of table	

Table 3.86: Relational functions which work on SYCL Host and device, are available in the `cl::sycl` namespace. They correspond to Table 6.14 of the OpenCL 1.2 specification [1]

3.10.8 Vector data and store functions

The functionality from the OpenCL functions as defined in in the OpenCL 1.2 specification document [1, par. 6.12.7] is available in SYCL through the `vec` class in section 3.7.2.

3.10.9 Synchronization Functions

In SYCL the OpenCL *synchronization* functions are available through the item class [3.5.1.4](#), as they are applied to work-item for local or global address spaces. Please see [3.55](#).

3.10.10 `printf` function

The functionality of the `printf` function is covered by the `cl::sycl::stream` class [3.9](#), which has the capability to print to standard output all the SYCL classes and primitives and covers the capabilities defined in the OpenCL 1.2 specification document [[1](#), par. 6.12.13].

4. SYCL Support of Non-Core OpenCL Features

OpenCL apart from *core* features that are supported in *every* platform, has *optional* features as well as *extensions* that are only supported in some platforms. The *optional* features, as described in the specification [1], and the OpenCL “*KHR*” *extensions*, as described in the extension specification [2], are supported by the SYCL framework, but the ability to use them is completely dependent on the underlying OpenCL platforms. A SYCL implementation may support some vendor extensions in order to enable optimizations on certain platforms.

All OpenCL extensions are available through SYCL interoperability with OpenCL C, so all the extensions can be used through the OpenCL API as described in the extensions specification [2].

When running command groups on the host device, not all extensions are required to be available. The extensions available for the *host* are available to query in the same way as for SYCL devices, see Table 3.5.

4.1 Enable extensions in a SYCL kernel

In order to enable extensions in an OpenCL kernel the following compiler directive is used:

```
#pragma OPENCL EXTENSION <extension_name> : <behaviour>
```

The keyword *extension_name* can be:

- **all**, which refers to all the extensions available on a platform
- an **extension name** from the available extensions on a platform.

The keyword *behaviour* can be:

- **enable**: it will enable the extension specified by *extension_name* if it is available on the platform or otherwise it triggers a compiler warning. If *all* is specified in the *extension_name* then it will enable all extensions available.
- **disable**: it will disable all or any extension provided in the *extension_name*.

The following table 4.1 summarizes the levels of SYCL support to the API extensions for OpenCL 1.2 [?]. These extensions can be supported using OpenCL/SYCL interoperability API or by the extended SYCL API calls. This only applies for using them in the framework and only for devices that are supporting these extensions.

Continued on next page

Table 4.1: SYCL support for OpenCL 1.2 API extensions.

Extension	Support using SYCL/OpenCL API	Support using SYCL API
-----------	-------------------------------	------------------------

Extension	Support using SYCL/OpenCL API	Support using SYCL API
cl_khr_int64_base_atomics	Yes	Yes
cl_khr_int64_extended_atomics	Yes	Yes
cl_khr_fp16	Yes	Yes
cl_khr_3d_image_writes	Yes	Yes
cl_khr_gl_sharing	Yes	Yes
cl_apple_gl_sharing	Yes	Yes
cl_khr_d3d10_sharing	Yes	No
cl_khr_d3d11_sharing	Yes	No
cl_khr_dx9_media_sharing	Yes	No
End of table		

Table 4.1: SYCL support for OpenCL 1.2 API extensions.

4.2 Half Precision Floating-Point

The half precision floating-point data scalar and vector types are supported in the SYCL system. The SYCL host device supports those types, however they are optional on an OpenCL device and the developer always needs to check whether the device the kernel is going to run on has the corresponding extension.

The extension name is `cl_khr_fp16` and it needs to be used in order to enable the usage of the half data type on an SYCL OpenCL device.

The half type class, along with any OpenCL macros and definitions, is defined in the namespace `cl::sycl` as `half`. The vector type of half is supported sizes 2, 3, 4, 8 and 16 using the SYCL vectors (§ 3.7.2) along with all the methods supported for vectors.

The conversion rules follows the same rules as in the OpenCL 1.2 extensions specification [2, par. 9.5.1].

The math, common, geometric and relational functions can take `cl::sycl::opencl::half` as a type as they are defined in [2, par. 9.5.2, 9.5.3, 9.5.4, 9.5.5]. The valid type for the functions defined for half is described by the generic type name *genhalf* is described in table 4.2.

Generic type name	Description
<code>genhalf</code>	<code>cl::sycl::half</code> , <code>cl::sycl::half2</code> , <code>cl::sycl::half3</code> , <code>cl::sycl::half4</code> , <code>cl::sycl::half8</code> , <code>cl::sycl::half16</code>
End of table	

Table 4.2: Generic type name description for all valid types of kernel function parameters. [1]

The elementary floating-point functions available for SYCL host and device is extended to allow half as input. If the half type is given as a parameter then the allowed error in ULP(Unit in the Last Place) is less than 8192. They correspond to Table 6.9 of the OpenCL 1.2 specification [1]

Math function	Description
<code>genhalf cos (genhalf x)</code>	Compute cosine. x must be in the range -216 to +216.
<code>genhalf divide (genhalf x, genhalf y)</code>	Compute x / y .
<code>genhalf exp (genhalf x)</code>	Compute the base- e exponential of x .
<code>genhalf exp2 (genhalf x)</code>	Compute the base- 2 exponential of x .
<code>genhalf exp10 (genhalf x)</code>	Compute the base- 10 exponential of x .
<code>genhalf log (genhalf x)</code>	Compute natural logarithm.
<code>genhalf log2 (genhalf x)</code>	Compute a base 2 logarithm.
<code>genhalf log10 (genhalf x)</code>	Compute a base 10 logarithm.
<code>genhalf powr (genhalf x, genhalf y)</code>	Compute x to the power y , where $x \geq 0$.
<code>genhalf recip (genhalf x)</code>	Compute reciprocal.
<code>genhalf rsqrt (genhalf x)</code>	Compute inverse square root.
<code>genhalf sin (genhalf x)</code>	Compute sine. x must be in the range -216 to +216.
<code>genhalf sqrt (genhalf x)</code>	Compute square root.
<code>genhalf tan (genhalf x)</code>	Compute tangent. x must be in the range -216 to +216.
End of table	

Table 4.3: Extended elementary functions which work on SYCL host and device.

4.3 Writing to 3D image memory objects

The `accessor` class for target `access::target::image` in SYCL support methods 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*.

4.4 Interoperability with OpenGL

OpenCL has a standard extension that allows interoperability with OpenGL objects. The features described in this section are only defined within SYCL if the underlying OpenCL implementation supports the OpenCL/OpenGL interoperability extension (`cl_khr_gl_sharing`).

4.4.1 OpenCL/OpenGL extensions to the context class

If the `cl_khr_gl_sharing` extension is present then the developer can create an OpenCL context from an OpenGL context by providing the corresponding attribute names and values to *properties* for the devices chosen by device selector. Table 3.8 has the additions shown on Table 4.5.

cl_context_properties flag	Description
CL_GL_CONTEXT_KHR	OpenGL context handle (default: 0)
CL_EGL_DISPLAY_KHR	CGL share group handle (default: 0)
CL_GLX_DISPLAY_KHR	EGLDisplay handle (default: EGL_NO_DISPLAY)
CL_WGL_HDC_KHR	X handle (default: None)
CL_CGL_SHAREGROUP_KHR	HDC handle (default: 0)
End of table	

Table 4.4: Additional optional properties for creating context for SYCL/OpenGL sharing.

The following table 4.5 describes the additional methods of the context class defined for the OpenCL/OpenGL interop that are also available for SYCL/OpenGL interop. If the OpenGL extensions are not available then their behavior is implementation defined.

Methods	Description
<code>device</code> <code>get_gl_current_device ()</code>	Returns the OpenGL enabled device in the current context.
<code>vector_class<device></code> <code>get_gl_context_devices ()</code>	Returns the OpenGL supported devices in this context.
End of table	

Table 4.5: Extended `context` class interface.

The SYCL extension for creating OpenCL context from an OpenGL context is based on the OpenCL extension specification and all the capabilities and restrictions are based on it and developers and implementers are advised to refer to [2, sec. 9.6].

4.4.2 Sharing SYCL/OpenGL memory objects

It is possible to share objects between SYCL and OpenGL, if the corresponding OpenCL platform extensions for these are available on available platforms. OpenCL memory objects based on OpenGL objects can only be created only if the OpenCL context is created from an OpenGL share group object or context. As the latter are OS specific, the OpenCL extensions are platform specific as well. In MacOS X the extension `cl_apple_gl_sharing` needs to be available for this functionality. If it is Windows/Linux/Unix, then the extension `cl_khr_gl_sharing` needs to be available. All the OpenGL objects within the shared group used for the creation of the context can be used apart from the default OpenGL objects.

Any of the buffers or images created through SYCL using the shared group objects for OpenGL are invalid if the corresponding OpenGL context is destroyed through usage of the OpenGL API. If buffers or images are used after the destruction of the corresponding OpenGL context then the behaviour of the system is undefined.

4.4.2.1 SYCL/OpenGL extensions to SYCL buffer

A SYCL *buffer* can be created from an OpenGL buffer object but the lifetime of the SYCL buffer is bound to the lifetime of the OpenCL context given in order to create the buffer. The GL buffer has to be created a priori using the OpenGL API, although it doesn't need to be initialized. If the OpenGL buffer object is destroyed or otherwise manipulated through the OpenGL API, before its usage through SYCL is completed, then the behaviour is undefined.

The functionality of the buffer and the accessor class is retained as for any other SYCL buffer defined in this system.

Constructor	Description
<code>template <typename T, int dimensions = 1> buffer(context &clGLContext, GLuint glBufferObj)</code>	Constructs a buffer from a OpenCL/OpenGL interop context and a <code>gl_buffer</code> object.
End of table	

Table 4.6: Extended constructors for the `buffer` class.

The extended methods of the buffer class, which have defined only when behavior when the OpenGL extensions are available on device, otherwise its undefined.

Method	Description
<code>cl_gl_object_type get_gl_info (GLuint glBufferObj)</code>	Returns the <code>cl_gl_object_type</code> of the underlying OpenGL buffer.
End of table	

Table 4.7: Extended `buffer` class interface

4.4.2.2 SYCL/OpenGL extensions to SYCL image

A SYCL *image* can be created from an OpenGL *buffer*, from an OpenGL *texture* or from an OpenGL *renderbuffer*. However, the lifetime of the SYCL image is bound to the lifetime of the OpenCL/OpenGL context given in order to create the image and the OpenGL object's lifetime. The GL buffer, texture or renderbuffer has to be created a priori via the OpenGL API, although it doesn't need to be initialized. If the OpenGL object is destroyed or otherwise manipulated through the OpenGL API before its usage through SYCL is completed, then the behaviour is undefined.

Constructor	Description
<code>template<int dimensions = 1> image(context &clGLContext, GLuint glBufferObj)</code>	Creates an 1-D Image from an OpenGL buffer object.
<code>template<int dimensions = 2> image(context &clGLContext, GLuint glRenderbufferObj)</code>	Create a 2-D image from an OpenGL renderbuffer object.
Continued on next page	

Table 4.8: Additional optional *image* class constructors.

Constructor	Description
<pre>template<int dimensions = 1> image(context &clGlContext, GLenum textureTarget , GLuint glTexture, GLint glMiplevel)</pre>	<p>Creates a 1-D image from an OpenGL texture object with given textureTarget and mipmap level. The textureTarget can be one of the following:</p> <ul style="list-style-type: none"> • GL_TEXTURE_1D • GL_TEXTURE_1D_ARRAY • GL_TEXTURE_BUFFER
<pre>template<int dimensions = 2> image(context &clGlContext, GLenum textureTarget , GLuint glTexture, GLint glMiplevel)</pre>	<p>Creates a 2-D image from an OpenGL texture object with given textureTarget and mipmap level. The textureTarget can be one of the following:</p> <ul style="list-style-type: none"> • GL_TEXTURE_2D • GL_TEXTURE_2D_ARRAY • GL_TEXTURE_CUBE_MAP_POSITIVE_X • GL_TEXTURE_CUBE_MAP_POSITIVE_Y • GL_TEXTURE_CUBE_MAP_POSITIVE_Z • GL_TEXTURE_CUBE_MAP_NEGATIVE_X • GL_TEXTURE_CUBE_MAP_NEGATIVE_Y • GL_TEXTURE_CUBE_MAP_NEGATIVE_Z • GL_TEXTURE_RECTANGLE
<pre>template<int dimensions = 3> image(context &clGlContext, GLenum textureTarget , GLuint glTexture, GLint glMiplevel)</pre>	<p>Creates a 3-D image from an OpenGL texture object with given textureTarget and mipmap level. The textureTarget can be one of the following:</p> <ul style="list-style-type: none"> • GL_TEXTURE_3D
End of table	

Table 4.8: Additional optional *image* class constructors.

Method	Description
GLenum get_gl_texture_target ()	Returns the OpenGL texture_target corresponding to the underlying texture which the context was created with.
GLint get_gl_mipmap_level ()	Returns the mipmap level of the underlying texture.
End of table	

Table 4.9: Additional optional *image* class method.

The *texture* provided has to be an OpenGL texture created through the OpenGL API and has to be a valid 1D, 2D, 3D texture or 1D array, 2D array texture or a cubemap, rectangle or buffer texture object. The format and the dimensions provided for the miplevel of the texture are used to create the OpenCL image object. The format of the OpenGL texture or renderbuffer object needs to match the format of the OpenCL image format. The compatible formats are specified in Table 9.4 of the OpenCL 1.2 extensions document [2, par. 9.7.3.1] and are also included in Table 4.10. If the texture or renderbuffer has a different format than the ones specified in 4.10, it is not guaranteed that the image created will be mapped to the the original texture.

OpenGL internal format	Corresponding OpenCL image format (channel order, channel data type)
GL_RGBA8	CL_RGBA, CL_UNORM_INT8, CL_BGRA, CL_UNORM_INT8
GL_RGBA, GL_UNSIGNED_INT_8_8_8_8_REV	CL_RGBA, CL_UNORM_INT8
GL_BGRA, GL_UNSIGNED_INT_8_8_8_8_REV	CL_BGRA, CL_UNORM_INT8
GL_RGBA16	CL_RGBA, CL_UNORM_INT16
GL_RGBA8I, GL_RGBA8I_EXT	CL_RGBA, CL_SIGNED_INT8
GL_RGBA16I, GL_RGBA16I_EXT	CL_RGBA, CL_SIGNED_INT16
GL_RGBA32I, GL_RGBA32I_EXT	CL_RGBA, CL_SIGNED_INT32
GL_RGBA8UI, GL_RGBA8UI_EXT	CL_RGBA, CL_UNSIGNED_INT8
GL_RGBA16UI, GL_RGBA16UI_EXT	CL_RGBA, CL_UNSIGNED_INT16
GL_RGBA32UI, GL_RGBA32UI_EXT	CL_RGBA, CL_UNSIGNED_INT32
GL_RGBA16F, GL_RGBA16F_ARB	CL_RGBA, CL_HALF_FLOAT
GL_RGBA32F, GL_RGBA32F_ARB	CL_RGBA, CL_FLOAT
End of table	

Table 4.10: Mapping of GL internal format to CL image format (reference: [2, table 9.4])

4.4.2.3 SYCL/OpenGL extensions to SYCL accessors

In order for SYCL to support the OpenCL/OpenGL interoperability, the classes for buffers and images have to be extended so that the OpenCL memory objects to be created from OpenGL objects. This extension, apart from restrictions on the creation and the life-time of the OpenCL objects, also requires that before the usage of any of these objects in an OpenCL *command.queue* an *acquire* command has to be enqueued first. In SYCL, the accessor classes make sure that the data are made available on a device. For this extension the SYCL kernel has to capture and through the accessor classes acquire any targets that are declared as interoperability targets.

The required extension for the accessor class are shown on Table 4.11.

Enumerator access::target values	Description
access::target::cl_gl_buffer	access <i>buffer</i> which is created from an OpenGL buffer
access::target::cl_gl_image	access an <i>image</i> or that is created from an OpenGL shared object
End of table	

Table 4.11: Enumerator description for access::target

The SYCL system is responsible for synchronizing the OpenCL and OpenGL objects in use inside a `command_group` when the SYCL API is used and given that all the accessors for the buffers and images are marked as the interoperability targets.

4.4.2.4 SYCL/OpenGL extensions to SYCL events

In the case where the extension `cl_khr_gl_event` is available on a platform, the functionality for creating synchronizing OpenCL events with OpenGL events is available and the `event` class is extended to include them.

A SYCL event can be constructed from an OpenGL sync object with the extensions to the event class shown on Table 4.12.

Constructor	Description
<code>event(context &clGlContext, GL_sync syncObj)</code>	Creates an event which enables waiting on events to also include synchronization with OpenGL objects that are shared using the OpenCL/OpenGL context
End of table	

Table 4.12: Additional optional class constructors for `event` class.

Method	Description
<code>GL_sync get_gl_info ()</code>	Returns <code>GL_sync</code> object.
End of table	

Table 4.13: Additional optional class method for `event` class.

The specification of the underlying OpenCL/OpenGL interoperability system for synchronizing OpenCL event with OpenGL sync objects is in the OpenCL extensions specification [2, sec. 9.8].

4.4.2.5 Extension for depth and depth-stencil images

The extension `cl_khr_depth_images` adds support for depth images and the extension `cl_khr_gl_depth_images` allows sharing between OpenCL depth images and OpenGL depth or depth-stencil textures. The SYCL system doesn't add any additional functionality towards this extension and follows the OpenCL 1.2 Specification [2, sec. 9.12] for depth and depth-stencil images extension. All the image class constructors and methods of the SYCL API as described in Table 3.23 and 3.24 on page 74 are extended to enable the use of the same API when this extension is present. The API is able to support the type `image2d_depth_t` and `image2d_array_depth_t`. The OpenCL C API defined in [2, sec. 9.12] can be used as well with all the rules that apply for SYCL/OpenCL C interoperability.

5. SYCL Device Compiler

This section specifies the requirements of the SYCL device compiler. Most features described in this section relate to underlying OpenCL 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 *shared source device compiler* and a *single-source device compiler*.

A SYCL shared source 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 shared-source compilation. 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 parses 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 needs to have a *globally-visible* name. In the case of a functor, the name can be the typename of the functor, as long as it is globally-visible. In the case where it isn't, a globally-visible name has to be provided, as template parameter to the kernel invoking interface, as described in 3.5.3. In C++11, lambdas¹ do not have a globally-visible name, so a globally-visible typename has to be provided in the kernel invoking

¹C++14 lambdas have the same naming rules as C++11 lambdas.

interface, as described in [3.5.3](#).

- The kernel name has to be a unique identifier in the program.

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 Language restrictions for kernels

The extracted SYCL kernels need to be compiled by an OpenCL online or offline compiler and be executed by the OpenCL 1.2 runtime. The extracted kernels need to be OpenCL 1.2 compliant kernels and as such there are certain restrictions that apply to them.

The following restrictions are applied to device functions and kernels:

- Structures containing pointers may be shared but the value of any pointer passed between SYCL devices or between the host and a SYCL device is undefined.
- Memory storage allocation is not allowed in kernels, all memory allocation for the device is done on host, using accessor classes. 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.
- No virtual methods are allowed to be called in a SYCL kernel or any functions called by the kernel.
- No function pointers are allowed to be called in a SYCL kernel or any functions called by the kernel.
- No class with a *vtable* can be used in a SYCL kernel or any code included in the kernel.
- RTTI is disabled inside kernels.
- Exception-handling cannot be used inside a SYCL kernel or any code called from the kernel.
- Recursion is not allowed in a SYCL kernel or any code called from the kernel.
- Global variables are not allowed to be used in kernel code.
- Non-const static member variables are not allowed to be used in kernel code.
- The rules for kernels apply to both the kernel functors themselves and all functions, operators, methods, 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 [3.10](#) to find functions that are specified to be usable in kernels.

5.4 Compilation of functions

The SYCL device compiler parses an entire C++ source file supplied by the user. This also includes C++ header files, using `#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.

In SYCL, kernels are invoked using a kernel invoke function (e.g. `parallel_for`). The kernel invoke functions are templated by their kernel parameter, which is a function object (either a functor or a lambda). The code inside the function object that is invoked as a kernel is called the “kernel function”. The “kernel function” must always return void. Any function called by the kernel function is compiled for device and called a “device function”. Recursively, any function called by a device function is itself compiled as a device function.

For example, this source code shows three functions and a kernel invoke with comments explaining which functions need to be compiled for device.

```
void f ()
{
    // function "f" is not compiled for device

    single_task<class kernel_name>([=] ()
    {
        // This code compiled for device
        g (); // this line forces "g" to be compiled for device
    });
}

void g ()
{
    // called from kernel, so "g" is compiled for device
}

void h ()
{
    // not called from a device function, so not compiled for device
}
```

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 is a function that matches at least the C++11 specification, plus any extensions from the C++14 specification.

5.5 Built-in scalar data types

In a SYCL device compiler, the standard C++ fundamental types, including `int`, `short`, `long`, `long long int` need to be configured so that the device definitions of those types match the host definitions of those types. 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.

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.

SYCL Scalar Datatypes	Description
<code>char</code>	a signed 8-bit integer.
<code>unsigned char</code>	an unsigned 8-bit integer.
<code>short int</code>	a signed integer of at least 16 bits and whose size must match the definition on the host.
<code>unsigned short int</code>	an unsigned integer of at least 16 bits and whose size must match the definition on the host.
<code>int</code>	a signed integer of at least 16 bits and whose size must match the definition on the host.
<code>unsigned int</code>	an unsigned integer of at least 16 bits and whose size must match the definition on the host.
<code>long int</code>	a signed integer of at least 32 bits and whose size must match the definition on the host.
<code>unsigned long int</code>	an unsigned integer of at least 32 bits and whose size must match the definition on the host.
<code>long long int</code>	an integer of at least 64 bits and whose size must match the definition on the host.
<code>unsigned long long int</code>	an unsigned integer of at least 64 bits and whose size must match the definition on the host.
<code>float</code>	a 32-bit IEEE 754 floating-point value.
<code>double</code>	a 64-bit IEEE 754 floating-point value.
<code>half</code>	a 16-bit IEEE 754-2008 half-precision floating-point value.
<code>size_t</code>	the unsigned integer type of the result of the <code>sizeof</code> operator on host.
End of table	

Table 5.1: SYCL compiler fundamental scalar datatypes

The SYCL device compiler also supports the OpenCL C scalar types which map to the OpenCL C language fundamental scalar datatypes [3.73](#).

5.6 Preprocessor directives and macros

The standard C++ preprocessing directives and macros are supported.

- `CL_SYCL_LANGUAGE_VERSION` substitutes an integer reflecting the version number of the SYCL language being supported by the device compiler. The version of SYCL defined in this document will have `CL_SYCL_LANGUAGE_VERSION` substitute the integer 120;

- `__FAST_RELAXED_MATH__` is used to determine if the `-cl-fast-relaxed-math` optimization option is specified in the build options given to the SYCL device compiler. This is an integer constant of 1 if the option is specified and undefined otherwise;
- `__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_TARGET_SPIR__` is defined to 1 if the source file is being compiled with a SYCL compiler which is producing OpenCL SPIR binary.
- `SYCL_EXTERNAL` is a macro which enables external linkage of SYCL functions and methods to be included in a SYCL kernel. For more details see [5.9.1](#)

5.7 Attributes

The `attribute` syntax defined in the OpenCL C specification is supported in SYCL.

The `vec_type_hint`, `work_group_size_hint` and `reqd_work_group_size` kernel attributes in OpenCL C apply to kernel functions, but this is not syntactically possible in SYCL. In SYCL, these attributes are legal on device functions and their specification is propagated down to any caller of those device functions, such that the kernel attributes are the sum of all the kernel attributes of all device functions called. If there are any conflicts between different kernel attributes, then the behaviour is undefined.

5.8 Address-space deduction

In SYCL, there are several different types of pointer, or reference:

- Accessors 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 OpenCL address space based on their access mode.
- Explicit pointer classes (e.g. `global_ptr`) contain an OpenCL address space. This allows the compiler to determine whether the pointer references global, local, constant or private memory.
- C++ pointer and reference types (e.g. `int*`) are allowed within SYCL kernels. They can be constructed from the address of local variables, from explicit pointer classes, or from accessors. In all cases, a SYCL device compiler will need to auto-deduce the address space.

Inside kernels, conversions between accessors to buffers, explicit pointer classes and C++ pointers are allowed as long as they reference the same datatype and have compatible qualifiers and address spaces.

If a kernel function or device function contains a pointer or reference type, then 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 have the

address space of the explicit pointer class.

- If a variable is declared as a pointer type, but initialized in its declaration to a pointer value with an already-deduced 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.
- 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. This default assumption is expected to change to be the generic address space for OpenCL versions that support the generic address space.

It is illegal to assign a pointer value of one address space to a pointer variable of a different address space.

5.9 SYCL offline linking

5.9.1 SYCL functions and methods linkage

The default behavior in SYCL applications is that all the definitions and declarations of the functions and methods 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 is implementation defined. It is the only requirement in the SYCL system for any function or method to be able to be linked against a SYCL application or library. The SYCL linkage mechanism is implementation defined, but the existence of the macro is required across all implementations.

5.9.2 Offline linking with OpenCL C libraries

The SYCL system supports external offline linking of OpenCL C libraries with a SYCL application. An OpenCL C function can be included and used in a SYCL program by defining it as an `extern "C"` function and adding the OpenCL library to the SYCL program. Any kernel which uses the external function needs to be included in a SYCL program which is linked against the OpenCL C library.

The data types for SYCL/OpenCL C interoperability are defined in [3.73](#). Only those data types can be used in the `extern "C"` declaration. These data types are invalid in an OpenCL C kernel, but these should be used in case of interop with a SYCL kernel and can be converted to and from the C++ fundamental types that are the default types in SYCL. The underlying OpenCL built-in types for pointers and vectors are defined as `typedefs` within the SYCL `vec` and explicit pointer types. The `vec` class contains a `vector_t typedef` for the underlying OpenCL C data type, while the explicit pointer classes contain a `pointer_t typedef` for the underlying OpenCL C pointer type.

A. Glossary

The purpose of this glossary is to define the key concepts involved in specifying OpenCL SYCL. This section includes definitions of terminology used throughout the specification document.

Accessor: An accessor is an interface which allows a kernel function to access data maintained by a buffer.

Application scope: The application scope is the normal C++ source code in the application, outside of command groups and kernels.

Buffer: A buffer is an interface which maintains an area of memory which is to be accessed by a kernel function. It represents storage of data only, with access to that data achieved via accessors. The storage is managed by the SYCL runtime, but may involve OpenCL buffers.

Barrier: SYCL barriers are the same as OpenCL barriers. In SYCL, OpenCL's command-queue-barriers are created automatically on demand by the SYCL runtime to ensure kernels are executed in a semantically-correct order across multiple OpenCL contexts and queues. OpenCL's work-group barriers are available as an intrinsic function (same as in OpenCL) or generated automatically by SYCL's hierarchical parallel-for loops.

Command Group Functor: All of the OpenCL commands, memory object creation, copying, mapping and synchronization operations to correctly execute a kernel on a device are defined in a functor and called a *command group functor*. Command group functors executed in different threads are added to queues atomically, so it is safe to *submit* command group functors operating on shared queues, buffers and images.

Command Group Scope: The command group scope is the scope defined by the command group functor.

Command Queue: SYCL's command queues abstract the OpenCL command queue functionality and add a SYCL-specific host command queue, which executes SYCL kernels on the host.

Constant Memory: “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.” As defined in [1, p.15]

Device: SYCL's devices encapsulate OpenCL devices and add support for executing SYCL kernels on host.

Device Compiler: A SYCL device compiler is a compiler that is capable of taking in C++ source code containing SYCL kernels and outputting a binary object suitable for executing on an OpenCL device.

Functor: Functors are a concept from C++. An alternative name for functions in C++ is “function object”. A functor is a C++ class with an **operator()** method that enables the object to be executed in a way that looks like a function call, but where the object itself is also passed in as a parameter.

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 an N-dimensional value that starts at (0, 0, ...0).

Global Memory: As in OpenCL, global memory is a memory region accessible to all work-items executing in a

context. Buffers are mapped or copied into global memory for individual contexts inside the SYCL runtime in order to enable accessors within command groups to give access to buffers from kernels.

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.

Group Range: A group range is the range specifying the size of the work group.

Host: As in OpenCL, the host is the system that executes the SYCL API and the rest of the application.

Host pointer: A pointer to memory that is in the virtual address space on the host.

ID: An id is a one, two or three dimensional vector of integers. There are several different types of ID in SYCL: global ID, local ID, group ID. These different IDs are used to define work items

Image: Images in SYCL, like buffers, are abstractions of the OpenCL equivalent. As in OpenCL, an image stores a two- or three-dimensional structured array. The SYCL runtime will map or copy images to OpenCL images in OpenCL contexts in order to execute semantically correct kernels in different OpenCL contexts. Images are also accessible on the host via the various SYCL accessors available.

Implementation defined: Behavior that is explicitly allowed to vary between conforming implementations of SYCL. A SYCL implementer is required to document the implementation defined behavior.

Item ID: An item id is an interface used to retrieve the global id, group id and local id of a work item.

Kernel: A SYCL kernel is a C++ functor or lambda function that is compiled to execute on a device. There are several ways to define SYCL kernels defined in the SYCL specification. It is also possible in SYCL to use OpenCL kernels as specified in the OpenCL specification. Kernels can execute on either an OpenCL device or on the host.

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.

Kernel Scope: The scope inside the kernel functor or lambda is called kernel scope. Also, any function or method 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.

Local ID: A local id is an id which specifies a work items location within a group range.

Local Memory: As in OpenCL, local memory is a memory region associated with a work-group and accessible only by work-items in that work-group.

NDRange: An NDRange consists of two vectors of integers of one, two or three-dimensions that define the total number of work items to execute as well as the size of the work groups that the work items are to be executed within.

Platform: A platform in SYCL is encapsulates an OpenCL platform as defined in the OpenCL specification.

Private Memory: As in OpenCL, private memory is a region of memory private to a work-item. Variables defined in one work-items private memory are not visible to another work-item.

Program Object: A program object in SYCL is an OpenCL program object encapsulated in A SYCL class. It contains OpenCL kernels and functions compiled to execute on OpenCL devices. A program object can be

generated from SYCL C++ kernels by the SYCL runtime, or obtained from an OpenCL implementation.

Shared Source Build System: A shared source build system means that a single source file passed through both a host compiler and one or more device compilers. This enables multiple devices, instruction sets and binary formats to be produced from the same source code and integrated into the same piece of software.

SYCL Runtime: A SYCL runtime is an implementation of the SYCL runtime specification. The SYCL runtime manages the different OpenCL platforms, devices, contexts as well as the mapping or copying of data between host and OpenCL contexts to enable semantically correct execution of SYCL kernels.

Work-Group: A work group is an OpenCL work group, defined in OpenCL as a collection of related work-items that execute on a single compute unit. The work-items in the group execute the same kernel and share local memory and work-group barriers.

Work-Item: A work item is an OpenCL work item, defined in OpenCL as one of 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 executions within the collection by its global ID and local ID.

B. Interface of SYCL Classes in Full

B.1 Explicit pointer classes

The explicit pointer classes `global_ptr`, `local_ptr`, `private_ptr` and `constant_ptr` are defined in 3.4.7.1. The available functions for these classes in full are the following:

```
namespace cl {
namespace sycl {

template <typename ElementType>
class global_ptr {
public:
    // Implementation defined pointer type that corresponds to the SYCL/OpenCL
    // interoperability type for OpenCL C functions
    typedef __undefined__ pointer_t;

    global_ptr(pointer_t); // global pointer

    template <access::mode Mode>
    global_ptr(accessor<ElementType, 1, Mode, global_buffer>);

    global_ptr(const global_ptr&);

    ~global_ptr();

    ElementType& operator*();
    ElementType& operator[](size_t i);

    // implementation defined implicit conversion
    // to OpenCL C pointer types.
    operator pointer_t();
};

template <typename ElementType>
class constant_ptr {
public:
    // Implementation defined pointer type that corresponds to the SYCL/OpenCL
    // interoperability type for OpenCL C functions
    typedef __undefined__ pointer_t;

    constant_ptr(pointer_t); // constant pointer

    template <access::mode Mode>
    constant_ptr(accessor<ElementType, 1, Mode, constant_buffer>);
};
```

```

ElementType& operator*();
ElementType& operator[](size_t i);

// implementation defined implicit conversion
// to OpenCL C pointer types.
operator pointer_t();
};

template <typename ElementType>
class local_ptr {
public:
// Implementation defined pointer type that corresponds to the SYCL/OpenCL
// interoperability type for OpenCL C functions
typedef __undefined__ pointer_t;

local_ptr(pointer_t); // local pointer

template <access::mode Mode>
local_ptr(accessor<ElementType, 1, Mode, local_buffer>);

ElementType& operator*();
ElementType& operator[](size_t i);

// implementation defined implicit conversion
// to OpenCL C pointer types.
operator pointer_t();
};

template <typename ElementType>
class private_ptr {
public:
// Implementation defined pointer type that corresponds to the SYCL/OpenCL
// interoperability type for OpenCL C functions
typedef __undefined__ pointer_t;

private_ptr(pointer_t); // private pointer

ElementType& operator*();
ElementType& operator[](size_t i);

// implementation defined implicit conversion
// to OpenCL C pointer types.
operator pointer_t();
};
} // namespace sycl
} // namespace cl

template <typename ElementType>
bool operator==(const global_ptr<ElementType>& lhs,
                const global_ptr<ElementType>& rhs);
template <typename ElementType>
bool operator!=(const global_ptr<ElementType>& lhs,
                const global_ptr<ElementType>& rhs);
template <typename ElementType>

```



```

bool operator<=(const private_ptr<ElementType>& lhs, nullptr_t rhs);
template <typename ElementType>
bool operator<=(nullptr_t lhs, const private_ptr<ElementType>& rhs);

```

B.2 Multi pointer pointer class

The `multi_ptr` class is defined in 3.4.7.2. The available functions for this class in full is the following:

```

namespace cl {
namespace sycl {
namespace access {
enum class address_space : int {
    global_space,
    local_space,
    constant_space,
    private_space
};
} // namespace access

template <typename ElementType, access::address_space Space>
class multi_ptr {
public :
    // Implementation defined pointer type that corresponds to the SYCL/OpenCL
    // interoperability type for OpenCL C functions
    typedef __undefined__ pointer_t;

    const address_space space;

    multi_ptr(pointer_t);
    multi_ptr(const multi_ptr&);
    ~multi_ptr();

    ElementType &operator*();
    ElementType &operator[](size_t i);

    // Only if Space == global_space
    operator global_ptr<ElementType>();
    global_ptr<ElementType> pointer();

    // Only if Space == local_space
    operator local_ptr<ElementType>();
    local_ptr<ElementType> pointer();

    // Only if Space == constant_space
    operator constant_ptr<ElementType>();
    constant_ptr<ElementType> pointer();

    // Only if Space == private_space
    operator private_ptr<ElementType>();
    private_ptr<ElementType> pointer();
};

```



```

template <typename ElementType, access::address_space Space>
multi_ptr<ElementType, Space> make_ptr(pointer_t);
} // namespace sycl
} // namespace cl

template <typename ElementType, access::address_space Space>
bool operator==(const multi_ptr<ElementType, Space>& lhs,
                const multi_ptr<ElementType, Space>& rhs);
template <typename ElementType, access::address_space Space>
bool operator!=(const multi_ptr<ElementType, Space>& lhs,
                const multi_ptr<ElementType, Space>& rhs);
template <typename ElementType, access::address_space Space>
bool operator<(const multi_ptr<ElementType, Space>& lhs,
               const multi_ptr<ElementType, Space>& rhs);
template <typename ElementType, access::address_space Space>
bool operator>(const multi_ptr<ElementType, Space>& lhs,
               const multi_ptr<ElementType, Space>& rhs);
template <typename ElementType, access::address_space Space>
bool operator<=(const multi_ptr<ElementType, Space>& lhs,
                const multi_ptr<ElementType, Space>& rhs);
template <typename ElementType, access::address_space Space>
bool operator>=(const multi_ptr<ElementType, Space>& lhs,
                const multi_ptr<ElementType, Space>& rhs);
template <typename ElementType, access::address_space Space>
bool operator!=(const multi_ptr<ElementType, Space>& lhs, nullptr_t rhs);
template <typename ElementType, access::address_space Space>
bool operator!=(nullptr_t lhs, const multi_ptr<ElementType, Space>& rhs);
template <typename ElementType, access::address_space Space>
bool operator==(const multi_ptr<ElementType, Space>& lhs, nullptr_t rhs);
template <typename ElementType, access::address_space Space>
bool operator==(nullptr_t lhs, const multi_ptr<ElementType, Space>& rhs);
template <typename ElementType, access::address_space Space>
bool operator>(const multi_ptr<ElementType, Space>& lhs, nullptr_t rhs);
template <typename ElementType, access::address_space Space>
bool operator>(nullptr_t lhs, const multi_ptr<ElementType, Space>& rhs);
template <typename ElementType, access::address_space Space>
bool operator<(const multi_ptr<ElementType, Space>& lhs, nullptr_t rhs);
template <typename ElementType, access::address_space Space>
bool operator<(nullptr_t lhs, const multi_ptr<ElementType, Space>& rhs);
template <typename ElementType, access::address_space Space>
bool operator>=(const multi_ptr<ElementType, Space>& lhs, nullptr_t rhs);
template <typename ElementType, access::address_space Space>
bool operator>=(nullptr_t lhs, const multi_ptr<ElementType, Space>& rhs);

```

B.3 range class

The `range` class is defined in [3.5.1.1](#). The available functions for this class in full is the following:

```

namespace cl {
namespace sycl {
template <size_t dimensions>
struct range {
    range(const range<dimensions> &);

    range(size_t x); // When dimensions==1
    range(size_t x, size_t y); // When dimensions==2
    range(size_t x, size_t y, size_t z); // When dimensions==3

    size_t get(int dimension) const;
    size_t &operator[](int dimension);

    range &operator=(const range &rhs);
    range &operator+=(const range &rhs);
    range &operator*=(const range &rhs);
    range &operator/=(const range &rhs);
    range &operator%=(const range &rhs);
    range &operator>>=(const range &rhs);
    range &operator<<=(const range &rhs);
    range &operator&=(const range &rhs);
    range &operator^=(const range &rhs);
    range &operator|=(const range &rhs);

    size_t size() const;
};

template <size_t dimensions>
bool operator==(const range<dimensions> &a, const range<dimensions> &b);
template <size_t dimensions>
bool operator!=(const range<dimensions> &a, const range<dimensions> &b);
template <size_t dimensions>
bool operator>(const range<dimensions> &a, const range<dimensions> &b);
template <size_t dimensions>
bool operator<(const range<dimensions> &a, const range<dimensions> &b);
template <size_t dimensions>
bool operator<=(const range<dimensions> &a, const range<dimensions> &b);
template <size_t dimensions>
bool operator>=(const range<dimensions> &a, const range<dimensions> &b);
template <size_t dimensions>
range<dimensions> operator*(const range<dimensions> &a,
                           const range<dimensions> &b);
template <size_t dimensions>
range<dimensions> operator/(const range<dimensions> &a,
                           const range<dimensions> &b);
template <size_t dimensions>
range<dimensions> operator+(const range<dimensions> &a,
                           const range<dimensions> &b);
template <size_t dimensions>
range<dimensions> operator-(const range<dimensions> &a,
                           const range<dimensions> &b);
template <size_t dimensions>
range<dimensions> operator%(const range<dimensions> &a,
                           const range<dimensions> &b);
template <size_t dimensions>

```


B.4 id class

The `id` class is defined in 3.5.1.3. The available functions for this class in full is the following:

```
namespace cl {
namespace sycl {
template <size_t dimensions>
struct id {
    id(size_t x); // When dimensions==1
    id(size_t x, size_t y); // When dimensions==2
    id(size_t x, size_t y, size_t z); // When dimensions==3
    id(const id<dimensions> & rhs);
    id(const range<dimensions> & rangeSize);
    id(const item<dimensions> & rhs);

    size_t get(int dimension) const;
    size_t &operator[](int dimension);
    operator size_t(); // When dimensions==1

    id &operator=(const id & rhs);
    id &operator+=(const id & rhs);
    id &operator*=(const id & rhs);
    id &operator/=(const id & rhs);
    id &operator%=(const id & rhs);
    id &operator>>=(const id & rhs);
    id &operator<<=(const id & rhs);
    id &operator&=(const id & rhs);
    id &operator^=(const id & rhs);
    id &operator|=(const id & rhs);
};

template <size_t dimensions>
bool operator==(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
bool operator!=(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
bool operator>(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
bool operator<(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
bool operator<=(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
bool operator>=(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator*(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator/(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator+(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator-(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator%(const id<dimensions> &a, const id<dimensions> &b);
```

```

template <size_t dimensions>
id<dimensions> operator<<(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator>>(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator&(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator|(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator^(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator&&(const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator|| (const id<dimensions> &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator*(const size_t &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator*(const id<dimensions> &a, const size_t &b);
template <size_t dimensions>
id<dimensions> operator/(const size_t &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator/(const id<dimensions> &a, const size_t &b);
template <size_t dimensions>
id<dimensions> operator+(const size_t &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator+(const id<dimensions> &a, const size_t &b);
template <size_t dimensions>
id<dimensions> operator-(const size_t &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator-(const id<dimensions> &a, const size_t &b);
template <size_t dimensions>
id<dimensions> operator%(const size_t &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator%(const id<dimensions> &a, const size_t &b);
template <size_t dimensions>
id<dimensions> operator<<(const size_t &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator<<(const id<dimensions> &a, const size_t &b);
template <size_t dimensions>
id<dimensions> operator>>(const size_t &a, const id<dimensions> &b);
template <size_t dimensions>
id<dimensions> operator>>(const id<dimensions> &a, const size_t &b);
} // namespace sycl
} // namespace cl

```

B.5 vec class

The vec class is defined in 3.7.2. The available functions for this class in full is the following:

```

namespace cl {
namespace sycl {

```

```

template <typename dataT, int numElements>
class vec {
public:
    typedef dataT element_type;
    //Underlying OpenCL type
    typedef __undefined__ vector_t;

    vec();

    explicit vec(const dataT &arg);

    vec(const T0 &arg0... args);

    vec(const vec<dataT, numElements> &rhs);

    size_t get_count();

    size_t get_size();

    template <typename asDataT, int width>
    vec<asDataT, width> as() const;

    // genvector is a generic typename for describing
    // all OpenCL/SYCL types.
    operator __genvector__() const;

    // arithmetic operators
    vec operator+(const vec &rhs) const;
    vec operator-(const vec &rhs) const;
    vec operator*(const vec &rhs) const;
    vec operator/(const vec &rhs) const;
    vec operator%(const vec &rhs) const;
    vec operator++(const vec &rhs) const;
    vec operator++();
    vec operator++(int);
    vec operator--();
    vec operator--(int);
    vec operator+(const dataT &rhs) const;
    vec operator-(const dataT &rhs) const;
    vec operator*(const dataT &rhs) const;
    vec operator/(const dataT &rhs) const;
    vec operator%(const dataT &rhs) const;

    // bitwise and logical operators
    vec operator|(const vec &rhs) const;
    vec operator|(const dataT &rhs) const;
    vec operator^(const vec &rhs) const;
    vec operator^(const dataT &rhs) const;
    vec operator&&(const vec<dataT, numElements> &rhs) const;
    vec operator&&(const dataT &rhs) const;
    vec operator|||(const vec<dataT, numElements> &rhs) const;
    vec operator|||(const dataT &rhs) const;
    vec operator>>(const vec<dataT, numElements> &rhs) const;
    vec operator>>(const dataT &rhs) const;
    vec operator<<(const vec &rhs) const;

```

```

vec operator<<(const dataT &rhs) const;
vec operator~();
vec operator!();

// assignment operators
vec operator+=(const vec &rhs);
vec operator+=(const dataT &rhs);
vec operator-=(const vec &rhs);
vec operator-=(const dataT &rhs);
vec operator*=(const vec &rhs);
vec operator*=(const dataT &rhs);
vec operator/=(const vec &rhs);
vec operator/=(const dataT &rhs);
vec operator|=(const vec &rhs);
vec operator|=(const dataT &rhs);
vec operator^=(const vec &rhs);
vec operator^=(const dataT &rhs);
vec operator<<=(const vec &rhs);
vec operator<<=(const dataT &rhs);
vec operator>>=(const vec &rhs);
vec operator>>=(const dataT &rhs);
vec operator&=(const vec &rhs);
vec operator&=(const dataT &rhs);
vec &operator=(const vec &rhs);
vec &operator=(const dataT &rhs);
vec &operator%=(const vec &rhs);
vec &operator%=(const dataT &rhs);

// relational operators
vec<int, numElements> operator==(const vec &rhs) const;
vec<int, numElements> operator!=(const vec &rhs) const;
vec<int, numElements> operator<=(const vec &rhs) const;
vec<int, numElements> operator>=(const vec &rhs) const;
vec<int, numElements> operator>(const vec &rhs) const;
vec<int, numElements> operator<(const vec &rhs) const;

// Swizzle methods (see notes)
swizzled_vec<T, out_dims> swizzle<elem s1, ...>();

#ifdef SYCL_SIMPLE_SWIZZLES
swizzled_vec<T, 4> xyzw();
...
#endif // #ifdef SYCL_SIMPLE_SWIZZLES
};
} // namespace sycl
} // namespace cl

```

C. Interface of Memory Object 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 3, for platform, context, device, and queue.

C.1 Platform Information Descriptors

The following interface includes all the information descriptors for the `platform` class as described in table 3.6.

```
namespace cl {
namespace sycl {
namespace info {
enum class platform : unsigned int {
    profile,
    version,
    name,
    vendor,
    extensions
};
}
```

C.2 Context Information Descriptors

The following interface includes all the information descriptors for the `context` class as described in table 3.9.

```
namespace cl {
namespace sycl {
namespace info {
bool gl_context_interop;
enum class context : int {
    reference_count,
    num_devices,
    gl_interop
};
} // info
} // sycl
} // cl
```


C.3 Device Information Descriptors

The following interface includes all the information descriptors for the `device` class as described in table 3.12.

```
namespace cl {
namespace sycl {
namespace info {
unsigned int device_fp_config;

unsigned int device_exec_capabilities;

unsigned int device_queue_properties;

enum class device_type : unsigned int {
    cpu,
    gpu,
    accelerator,
    custom,
    defaults,
    host,
    all
};

enum class device : int {
    device_type,
    vendor_id,
    max_compute_units,
    max_work_item_dimensions,
    max_work_item_sizes,
    max_work_group_size,
    preferred_vector_width_char,
    preferred_vector_width_short,
    preferred_vector_width_int,
    preferred_vector_width_long_long,
    preferred_vector_width_float,
    preferred_vector_width_double,
    preferred_vector_width_half,
    native_vector_width_char,
    native_vector_width_short,
    native_vector_width_int,
    native_vector_width_long_long,
    native_vector_width_float,
    native_vector_width_double,
    native_vector_width_half,
    max_clock_frequency,
    address_bits,
    max_mem_alloc_size,
    image_support,
    max_read_image_args,
    max_write_image_args,
    image2d_max_height,
    image2d_max_width,
    image3d_max_height,
```

```

image3d_max_widht,
image3d_mas_depth,
image_max_buffer_size,
image_max_array_size,
max_samplers,
max_parameter_size,
mem_base_addr_align,
single_fp_config,
double_fp_config,
global_mem_cache_type,
global_mem_cache_line_size,
global_mem_cache_size,
global_mem_size,
max_constant_buffer_size,
max_constant_args,
local_mem_type,
local_mem_size,
error_correction_support,
host_unified_memory,
profiling_timer_resolution,
endian_little,
is_available,
is_compiler_available,
is_linker_available,
execution_capabilities,
queue_properties,
built_in_kernels,
platform,
name,
vendor,
driver_version,
profile,
device_version,
opengl_version,
extensions,
printf_buffer_size,
preferred_interop_user_sync,
parent_device,
partition_max_sub_devices,
partition_properties,
partition_affinity_domain,
partition_type,
reference_count
};

enum class device_partition_property : int {
    unsupported,
    partition_equally,
    partition_by_counts,
    partition_by_affinity_domain,
    partition_affinity_domain_next_partitionable
};

enum class device_affinity_domain : int {
    unsupported,

```

```

    numa,
    L4_cache,
    L3_cache,
    L2_cache,
    next_partitionable
};

enum class device_partition_type : int {
    no_partition,
    numa,
    L4_cache,
    L3_cache,
    L2_cache,
    L1_cache
};

enum class local_mem_type : int {
    none,
    local,
    global
};

enum class fp_config : int {
    denorm,
    inf_nan,
    round_to_nearest,
    round_to_zero,
    round_to_inf,
    fma,
    correctly_rounded_divide_sqrt,
    soft_float
};

enum class global_mem_cache_type : int {
    none,
    read_only,
    write_only
};

enum class device_execution_capabilities : unsigned int {
    exec_kernel,
    exec_native_kernel
};
} // namespace info
} // namespace sycl
} // namespace cl

```

C.4 Queue Information Descriptors

The following interface includes all the information descriptors for the `queue` class as described in table 3.15.

```

namespace cl {
namespace sycl {
namespace info {
bool queue_profiling;
enum class queue : int {
    context,
    device,
    reference_count,
    properties
};
} // namespace info
} // namespace sycl
} // namespace cl

```

C.5 Kernel Information Descriptors

The following interface includes all the information descriptors for the [kernel](#) class as described in table 3.62.

```

namespace cl {
namespace sycl {
namespace info {
enum class kernel: int {
    function_name,
    num_args,
    reference_count,
    attributes
};
} // namespace info
} // namespace sycl
} // namespace cl

```

C.6 Program Information Descriptors

The following interface includes all the information descriptors for the [program](#) class as described in table 3.65.

```

namespace cl {
namespace sycl {
namespace info {
enum class program: int {
    reference_count,
    context,
    devices
};
} // namespace info
} // namespace sycl
} // namespace cl

```

C.7 Event Information Descriptors

The following interface includes all the information descriptors for the `event` class as described in table 3.62 and table 3.19.

```
namespace cl {
namespace sycl {
namespace info {
enum class event: int {
    command_type,
    command_execution_status,
    reference_count
};

enum class event_profiling : int {
    command_queued,
    command_submit,
    command_start,
    command_end
};
} // namespace info
} // namespace sycl
} // namespace cl
```

References

- [1] Khronos OpenCL Working Group, *The OpenCL Specification, version 1.2.19*, 2012. [Online]. Available: <https://www.khronos.org/registry/cl/specs/openc1-2.0.pdf>
- [2] —, *The OpenCL Extension Specification, version 1.2.19*, 2012. [Online]. Available: <http://www.khronos.org/registry/cl/specs/openc1-1.2-extensions.pdf>