

The **OpenCL** Specification

Version: 1.2

Document Revision: 19

Khronos OpenCL Working Group

Editor: Aaftab Munshi

| | |
|---|-----------|
| 1. INTRODUCTION | 12 |
| 2. GLOSSARY | 14 |
| 2.1 OpenCL Class Diagram..... | 20 |
| 3. THE OPENCL ARCHITECTURE..... | 22 |
| 3.1 Platform Model..... | 22 |
| 3.1.1 Platform Mixed Version Support | 23 |
| 3.2 Execution Model..... | 24 |
| 3.2.1 Execution Model: Context and Command Queues | 26 |
| 3.2.2 Execution Model: Categories of Kernels | 27 |
| 3.3 Memory Model | 27 |
| 3.3.1 Memory Consistency | 29 |
| 3.4 Programming Model..... | 29 |
| 3.4.1 Data Parallel Programming Model | 29 |
| 3.4.2 Task Parallel Programming Model | 30 |
| 3.4.3 Synchronization | 30 |
| 3.5 Memory Objects..... | 31 |
| 3.6 The OpenCL Framework..... | 31 |
| 4. THE OPENCL PLATFORM LAYER | 33 |
| 4.1 Querying Platform Info | 33 |
| 4.2 Querying Devices..... | 35 |
| 4.3 Partitioning a Device..... | 50 |
| 4.4 Contexts..... | 55 |
| 5. THE OPENCL RUNTIME | 62 |
| 5.1 Command Queues | 62 |
| 5.2 Buffer Objects..... | 67 |
| 5.2.1 Creating Buffer Objects | 67 |
| 5.2.2 Reading, Writing and Copying Buffer Objects | 73 |
| 5.2.3 Filling Buffer Objects | 85 |
| 5.2.4 Mapping Buffer Objects..... | 87 |
| 5.3 Image Objects | 91 |

| | | |
|-------------|--|------------|
| 5.3.1 | Creating Image Objects | 91 |
| 5.3.1.1 | Image Format Descriptor | 93 |
| 5.3.1.2 | Image Descriptor | 96 |
| 5.3.2 | Querying List of Supported Image Formats | 97 |
| 5.3.2.1 | Minimum List of Supported Image Formats | 99 |
| 5.3.3 | Reading, Writing and Copying Image Objects | 99 |
| 5.3.4 | Filling Image Objects | 106 |
| 5.3.5 | Copying between Image and Buffer Objects | 108 |
| 5.3.6 | Mapping Image Objects | 113 |
| 5.3.7 | Image Object Queries | 116 |
| 5.4 | Querying, Unmapping, Migrating, Retaining and Releasing Memory Objects..... | 119 |
| 5.4.1 | Retaining and Releasing Memory Objects..... | 119 |
| 5.4.2 | Unmapping Mapped Memory Objects..... | 121 |
| 5.4.3 | Accessing mapped regions of a memory object..... | 123 |
| 5.4.4 | Migrating Memory Objects..... | 124 |
| 5.4.5 | Memory Object Queries..... | 126 |
| 5.5 | Sampler Objects | 129 |
| 5.5.1 | Creating Sampler Objects | 129 |
| 5.5.2 | Sampler Object Queries | 131 |
| 5.6 | Program Objects | 133 |
| 5.6.1 | Creating Program Objects | 133 |
| 5.6.2 | Building Program Executables..... | 138 |
| 5.6.3 | Separate Compilation and Linking of Programs..... | 140 |
| 5.6.4 | Compiler Options | 146 |
| 5.6.4.1 | Preprocessor options | 146 |
| 5.6.4.2 | Math Intrinsic Options | 146 |
| 5.6.4.3 | Optimization Options | 147 |
| 5.6.4.4 | Options to Request or Suppress Warnings..... | 148 |
| 5.6.4.5 | Options Controlling the OpenCL C version..... | 148 |
| 5.6.4.6 | Options for Querying Kernel Argument Information | 149 |
| 5.6.5 | Linker Options | 149 |
| 5.6.5.1 | Library Linking Options..... | 149 |
| 5.6.5.2 | Program Linking Options..... | 150 |
| 5.6.6 | Unloading the OpenCL Compiler | 150 |
| 5.6.7 | Program Object Queries..... | 151 |
| 5.7 | Kernel Objects | 158 |
| 5.7.1 | Creating Kernel Objects..... | 158 |
| 5.7.2 | Setting Kernel Arguments..... | 161 |
| 5.7.3 | Kernel Object Queries..... | 163 |
| 5.8 | Executing Kernels | 171 |
| 5.9 | Event Objects..... | 179 |
| 5.10 | Markers, Barriers and Waiting for Events..... | 188 |
| 5.11 | Out-of-order Execution of Kernels and Memory Object Commands..... | 191 |
| 5.12 | Profiling Operations on Memory Objects and Kernels..... | 192 |
| 5.13 | Flush and Finish | 195 |

| | | |
|-------------|--|------------|
| 6. | THE OPENCL C PROGRAMMING LANGUAGE | 197 |
| 6.1 | Supported Data Types | 197 |
| 6.1.1 | Built-in Scalar Data Types | 197 |
| 6.1.1.1 | The half data type | 199 |
| 6.1.2 | Built-in Vector Data Types | 200 |
| 6.1.3 | Other Built-in Data Types | 201 |
| 6.1.4 | Reserved Data Types | 202 |
| 6.1.5 | Alignment of Types | 203 |
| 6.1.6 | Vector Literals | 203 |
| 6.1.7 | Vector Components | 204 |
| 6.1.8 | Aliasing Rules | 208 |
| 6.1.9 | Keywords | 208 |
| 6.2 | Conversions and Type Casting | 209 |
| 6.2.1 | Implicit Conversions | 209 |
| 6.2.2 | Explicit Casts | 209 |
| 6.2.3 | Explicit Conversions | 210 |
| 6.2.3.1 | Data Types | 211 |
| 6.2.3.2 | Rounding Modes | 211 |
| 6.2.3.3 | Out-of-Range Behavior and Saturated Conversions | 212 |
| 6.2.3.4 | Explicit Conversion Examples | 212 |
| 6.2.4 | Reinterpreting Data As Another Type | 213 |
| 6.2.4.1 | Reinterpreting Types Using Unions | 213 |
| 6.2.4.2 | Reinterpreting Types Using <code>as_type()</code> and <code>as_type<i>n</i>()</code> | 214 |
| 6.2.5 | Pointer Casting | 215 |
| 6.2.6 | Usual Arithmetic Conversions | 216 |
| 6.3 | Operators | 217 |
| 6.4 | Vector Operations | 223 |
| 6.5 | Address Space Qualifiers | 224 |
| 6.5.1 | <code>__global</code> (or <code>global</code>) | 225 |
| 6.5.2 | <code>__local</code> (or <code>local</code>) | 225 |
| 6.5.3 | <code>__constant</code> (or <code>constant</code>) | 226 |
| 6.5.4 | <code>__private</code> (or <code>private</code>) | 227 |
| 6.6 | Access Qualifiers | 228 |
| 6.7 | Function Qualifiers | 229 |
| 6.7.1 | <code>__kernel</code> (or <code>kernel</code>) | 229 |
| 6.7.2 | Optional Attribute Qualifiers | 229 |
| 6.8 | Storage-Class Specifiers | 232 |
| 6.9 | Restrictions | 233 |
| 6.10 | Preprocessor Directives and Macros | 236 |
| 6.11 | Attribute Qualifiers | 238 |
| 6.11.1 | Specifying Attributes of Types | 239 |
| 6.11.2 | Specifying Attributes of Functions | 241 |
| 6.11.3 | Specifying Attributes of Variables | 241 |
| 6.11.4 | Specifying Attributes of Blocks and Control-Flow-Statements | 243 |

| | | |
|-------------|--|------------|
| 6.11.5 | Extending Attribute Qualifiers | 243 |
| 6.12 | Built-in Functions..... | 244 |
| 6.12.1 | Work-Item Functions | 244 |
| 6.12.2 | Math Functions | 246 |
| 6.12.2.1 | Floating-point macros and pragmas | 254 |
| 6.12.3 | Integer Functions..... | 258 |
| 6.12.4 | Common Functions | 262 |
| 6.12.5 | Geometric Functions | 264 |
| 6.12.6 | Relational Functions | 266 |
| 6.12.7 | Vector Data Load and Store Functions | 269 |
| 6.12.8 | Synchronization Functions..... | 276 |
| 6.12.9 | Explicit Memory Fence Functions | 277 |
| 6.12.10 | Async Copies from Global to Local Memory, Local to Global Memory, and Prefetch..... | 278 |
| 6.12.11 | Atomic Functions..... | 281 |
| 6.12.12 | Miscellaneous Vector Functions | 284 |
| 6.12.13 | printf..... | 286 |
| 6.12.13.1 | printf output synchronization | 286 |
| 6.12.13.2 | printf format string | 286 |
| 6.12.13.3 | Differences between OpenCL C and C99 printf | 292 |
| 6.12.14 | Image Read and Write Functions..... | 294 |
| 6.12.14.1 | Samplers | 294 |
| 6.12.14.2 | Built-in Image Read Functions | 297 |
| 6.12.14.3 | Built-in Image Sampler-less Read Functions..... | 305 |
| 6.12.14.4 | Built-in Image Write Functions..... | 311 |
| 6.12.14.5 | Built-in Image Query Functions..... | 314 |
| 6.12.14.6 | Mapping image channels to color values returned by read_image and color values passed to write_image to image channels | 316 |
| 7. | OPENCL NUMERICAL COMPLIANCE | 318 |
| 7.1 | Rounding Modes..... | 318 |
| 7.2 | INF, NaN and Denormalized Numbers | 318 |
| 7.3 | Floating-Point Exceptions | 319 |
| 7.4 | Relative Error as ULPs..... | 319 |
| 7.5 | Edge Case Behavior | 324 |
| 7.5.1 | Additional Requirements Beyond C99 TC2 | 324 |
| 7.5.2 | Changes to C99 TC2 Behavior | 327 |
| 7.5.3 | Edge Case Behavior in Flush To Zero Mode..... | 328 |
| 8. | IMAGE ADDRESSING AND FILTERING | 329 |
| 8.1 | Image Coordinates | 329 |
| 8.2 | Addressing and Filter Modes | 329 |
| 8.3 | Conversion Rules..... | 336 |
| 8.3.1 | Conversion rules for normalized integer channel data types | 336 |
| 8.3.1.1 | Converting normalized integer channel data types to floating-point values..... | 336 |
| 8.3.1.2 | Converting floating-point values to normalized integer channel data types..... | 337 |

| | | |
|--|--|------------|
| 8.3.2 | Conversion rules for half precision floating-point channel data type | 339 |
| 8.3.3 | Conversion rules for floating-point channel data type | 339 |
| 8.3.4 | Conversion rules for signed and unsigned 8-bit, 16-bit and 32-bit integer channel data types | 340 |
| 8.4 | Selecting an Image from an Image Array | 341 |
| 9. | OPTIONAL EXTENSIONS | 342 |
| 10. | OPENCL EMBEDDED PROFILE | 343 |
| 11. | REFERENCES..... | 358 |
| APPENDIX A | | 359 |
| A.1 | Shared OpenCL Objects..... | 359 |
| A.2 | Multiple Host Threads | 360 |
| APPENDIX B — PORTABILITY | | 361 |
| APPENDIX C — APPLICATION DATA TYPES | | 366 |
| C.1 | Shared Application Scalar Data Types | 366 |
| C.2 | Supported Application Vector Data Types | 366 |
| C.3 | Alignment of Application Data Types | 367 |
| C.4 | Vector Literals | 367 |
| C.5 | Vector Components..... | 367 |
| C.5.1 | Named vector components notation | 368 |
| C.5.2 | High/Low vector component notation..... | 368 |
| C.5.3 | Native vector type notation | 369 |
| C.6 | Implicit Conversions | 369 |
| C.7 | Explicit Casts | 369 |
| C.8 | Other operators and functions | 370 |
| C.9 | Application constant definitions | 370 |
| APPENDIX D — OPENCL C++ WRAPPER API | | 372 |
| APPENDIX E — CL_MEM_COPY_OVERLAP | | 373 |

| | |
|---|------------|
| APPENDIX F – CHANGES | 376 |
| F.1 Summary of changes from OpenCL 1.0 | 376 |
| F.2 Summary of changes from OpenCL 1.1 | 378 |
| INDEX - APIS | 380 |

Copyright (c) 2008-2011 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 re-formatted 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 web-site 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, StreamInput, WebGL, COLLADA, 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

The OpenCL specification is the result of the contributions of many people, representing a cross section of the desktop, hand-held, and embedded computer industry. Following is a partial list of the contributors, including the company that they represented at the time of their contribution:

David Neto, Altera
Benedict Gaster, AMD
Bill Licea Kane, AMD
Ed Buckingham, AMD
Jan Civlin, AMD
Laurent Morichetti, AMD
Mark Fowler, AMD
Michael Houston, AMD
Michael Mantor, AMD
Norm Rubin, AMD
Ofer Rosenberg, AMD
Victor Odintsov, AMD
Aaftab Munshi, Apple
Alexandre Namaan, Apple
Anna Tikhonova, Apple
Abe Stephens, Apple
Bob Beretta, Apple
Chendi Zhang, Apple
David Hayward, Apple
Geoff Stahl, Apple
Giridhar Murthy, Apple
Ian Ollmann, Apple
Inam Rahman, Apple
James Shearer, Apple
Jeremy Sandmel, Apple
John Stauffer, Apple
Kathleen Danielson, Apple
Michael Larson, Apple
MonPing Wang, Apple
Nate Begeman, Apple
Tanya Lattner, Apple
Travis Brown, Apple
Anton Lokhmotov, ARM
Dave Shreiner, ARM
Hedley Francis, ARM
Robert Elliott, ARM
Scott Moyers, ARM
Tom Olson, ARM

Alastair Donaldson, Codeplay
Andrew Richards, Codeplay
Stephen Frye, Electronic Arts
Eric Schenk, Electronic Arts
Brian Murray, Freescale
Brian Horton, IBM
Brian Watt, IBM
Dan Brokenshire, IBM
Gordon Fossum, IBM
Greg Bellows, IBM
Joaquin Madruga, IBM
Mark Nutter, IBM
Joe Molleson, Imagination Technologies
Jon Parr, Imagination Technologies
Robert Quill, Imagination Technologies
James McCarthy, Imagination Technologies
Aaron Lefohn, Intel
Adam Lake, Intel
Andrew Brownsword, Intel
Andrew Lauritzen, Intel
Craig Kolb, Intel
Geoff Berry, Intel
John Kessenich, Intel
Josh Fryman, Intel
Hong Jiang, Intel
Larry Seiler, Intel
Matt Pharr, Intel
Michael McCool, Intel
Murali Sundaresan, Intel
Paul Lalonde, Intel
Stefanus Du Toit, Intel
Stephen Junkins, Intel
Tim Foley, Intel
Timothy Mattson, Intel
Yariv Aridor, Intel
Bill Bush, Kestrel Institute
Lindsay Errington, Kestrel Institute
Jon Leech, Khronos
Benjamin Bergen, Los Alamos National Laboratory
Marcus Daniels, Los Alamos National Laboratory
Michael Bourges Sevenier, Motorola
Jyrki Leskelä, Nokia
Jari Nikara, Nokia
Amit Rao, NVIDIA

Ashish Srivastava, NVIDIA
Bastiaan Aarts, NVIDIA
Chris Cameron, NVIDIA
Christopher Lamb, NVIDIA
Dibyapran Sanyal, NVIDIA
Guatam Chakrabarti, NVIDIA
Ian Buck, NVIDIA
Jason Sanders, NVIDIA
Jaydeep Marathe, NVIDIA
Jian-Zhong Wang, NVIDIA
Karthik Raghavan Ravi, NVIDIA
Kedar Patil, NVIDIA
Manjunath Kudlur, NVIDIA
Mark Harris, NVIDIA
Michael Gold, NVIDIA
Neil Trevett, NVIDIA
Rahul Joshi, NVIDIA
Richard Johnson, NVIDIA
Sean Lee, NVIDIA
Tushar Kashalikar, NVIDIA
Vinod Grover, NVIDIA
Xiangyun Kong, NVIDIA
Yogesh Kini, NVIDIA
Yuan Lin, NVIDIA
Alex Bourd, QUALCOMM
Andrzej Mamona, QUALCOMM
Chihong Zhang, QUALCOMM
David Garcia, QUALCOMM
David Ligon, QUALCOMM
Robert Simpson, QUALCOMM
Yanjun Zhang, S3 Graphics
Tasneem Brutch, Samsung
Thierry Lepley, STMicroelectronics
Alan Ward, Texas Instruments
Madhukar Budagavi, Texas Instruments
Brian Hutsell Vivante
Mike Cai, Vivante
Sumeet Kumar, Vivante
Henry Styles, Xilinx

1. Introduction

Modern processor architectures have embraced parallelism as an important pathway to increased performance. Facing technical challenges with higher clock speeds in a fixed power envelope, Central Processing Units (CPUs) now improve performance by adding multiple cores. Graphics Processing Units (GPUs) have also evolved from fixed function rendering devices into programmable parallel processors. As today's computer systems often include highly parallel CPUs, GPUs and other types of processors, it is important to enable software developers to take full advantage of these heterogeneous processing platforms.

Creating applications for heterogeneous parallel processing platforms is challenging as traditional programming approaches for multi-core CPUs and GPUs are very different. CPU-based parallel programming models are typically based on standards but usually assume a shared address space and do not encompass vector operations. General purpose GPU programming models address complex memory hierarchies and vector operations but are traditionally platform-, vendor- or hardware-specific. These limitations make it difficult for a developer to access the compute power of heterogeneous CPUs, GPUs and other types of processors from a single, multi-platform source code base. More than ever, there is a need to enable software developers to effectively take full advantage of heterogeneous processing platforms – from high performance compute servers, through desktop computer systems to handheld devices - that include a diverse mix of parallel CPUs, GPUs and other processors such as DSPs and the Cell/B.E. processor.

OpenCL (Open Computing Language) is an open royalty-free standard for general purpose parallel programming across CPUs, GPUs and other processors, giving software developers portable and efficient access to the power of these heterogeneous processing platforms.

OpenCL supports a wide range of applications, ranging from embedded and consumer software to HPC solutions, through a low-level, high-performance, portable abstraction. By creating an efficient, close-to-the-metal programming interface, OpenCL will form the foundation layer of a parallel computing ecosystem of platform-independent tools, middleware and applications. OpenCL is particularly suited to play an increasingly significant role in emerging interactive graphics applications that combine general parallel compute algorithms with graphics rendering pipelines.

OpenCL consists of an API for coordinating parallel computation across heterogeneous processors; and a cross-platform programming language with a well-specified computation environment. The OpenCL standard:

- ✚ Supports both data- and task-based parallel programming models
- ✚ Utilizes a subset of ISO C99 with extensions for parallelism
- ✚ Defines consistent numerical requirements based on IEEE 754
- ✚ Defines a configuration profile for handheld and embedded devices
- ✚ Efficiently interoperates with OpenGL, OpenGL ES and other graphics APIs

This document begins with an overview of basic concepts and the architecture of OpenCL, followed by a detailed description of its execution model, memory model and synchronization support. It then discusses the OpenCL platform and runtime API and is followed by a detailed description of the OpenCL C programming language. Some examples are given that describe sample compute use-cases and how they would be written in OpenCL. The specification is divided into a core specification that any OpenCL compliant implementation must support; a handheld/embedded profile which relaxes the OpenCL compliance requirements for handheld and embedded devices; and a set of optional extensions that are likely to move into the core specification in later revisions of the OpenCL specification.

2. Glossary

Application: The combination of the program running on the *host* and *OpenCL devices*.

Blocking and Non-Blocking Enqueue API calls: A *non-blocking enqueue API call* places a *command* on a *command-queue* and returns immediately to the host. The *blocking-mode enqueue API calls* do not return to the host until the command has completed.

Barrier: There are two types of *barriers* – a *command-queue barrier* and a *work-group barrier*.

- ✚ The OpenCL API provides a function to enqueue a *command-queue barrier* command. This *barrier* command ensures that all previously enqueued commands to a *command-queue* have finished execution before any following *commands* enqueued in the *command-queue* can begin execution.
- ✚ The OpenCL C programming language provides a built-in *work-group barrier* function. This *barrier* built-in function can be used by a *kernel* executing on a *device* to perform synchronization between *work-items* in a *work-group* executing the *kernel*. All the *work-items* of a *work-group* must execute the *barrier* construct before any are allowed to continue execution beyond the *barrier*.

Buffer Object: A memory object that stores a linear collection of bytes. Buffer objects are accessible using a pointer in a *kernel* executing on a *device*. Buffer objects can be manipulated by the host using OpenCL API calls. A *buffer object* encapsulates the following information:

- ✚ Size in bytes.
- ✚ Properties that describe usage information and which region to allocate from.
- ✚ Buffer data.

Built-in Kernel: A *built-in kernel* is a *kernel* that is executed on an OpenCL *device* or *custom device* by fixed-function hardware or in firmware. *Applications* can query the *built-in kernels* supported by a *device* or *custom device*. A *program object* can only contain *kernels* written in OpenCL C or *built-in kernels* but not both. See also *Kernel* and *Program*.

Command: The OpenCL operations that are submitted to a *command-queue* for execution. For example, OpenCL commands issue *kernels* for execution on a compute device, manipulate memory objects, etc.

Command-queue: An object that holds *commands* that will be executed on a specific *device*. The *command-queue* is created on a specific *device* in a *context*. *Commands* to a *command-queue* are queued in-order but may be executed in-order or out-of-order. Refer to *In-order Execution* and *Out-of-order Execution*.

Command-queue Barrier. See *Barrier*.

Compute Device Memory: This refers to one or more memories attached to the compute device.

Compute Unit: An OpenCL *device* has one or more *compute units*. A *work-group* executes on a single *compute unit*. A *compute unit* is composed of one or more *processing elements* and *local memory*. A *compute unit* may also include dedicated texture filter units that can be accessed by its processing elements.

Concurrency: A property of a system in which a set of tasks in a system can remain active and make progress at the same time. To utilize concurrent execution when running a program, a programmer must identify the concurrency in their problem, expose it within the source code, and then exploit it using a notation that supports concurrency.

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*.

Context: The environment within which the *kernels* execute and the domain in which synchronization and memory management is defined. The *context* includes a set of *devices*, the memory accessible to those *devices*, the corresponding memory properties and one or more *command-queues* used to schedule execution of a *kernel(s)* or operations on *memory objects*.

Custom Device: An OpenCL *device* that fully implements the OpenCL Runtime but does not support *programs* written in OpenCL C. A custom device may be specialized non-programmable hardware that is very power efficient and performant for directed tasks or hardware with limited programmable capabilities such as specialized DSPs. Custom devices are not OpenCL conformant. Custom devices may support an online compiler. Programs for custom devices can be created using the OpenCL runtime APIs that allow OpenCL programs to be created from source (if an online compiler is supported) and/or binary, or from *built-in kernels* supported by the *device*. See also *Device*.

Data Parallel Programming Model: Traditionally, this term refers to a programming model where concurrency is expressed as instructions from a single program applied to multiple elements within a set of data structures. The term has been generalized in OpenCL to refer to a model wherein a set of instructions from a single program are applied concurrently to each point within an abstract domain of indices.

Device: A *device* is a collection of *compute units*. A *command-queue* is used to queue *commands* to a *device*. Examples of *commands* include executing *kernels*, or reading and writing *memory objects*. OpenCL devices typically correspond to a GPU, a multi-core CPU, and other processors such as DSPs and the Cell/B.E. processor.

Event Object: An *event object* encapsulates the status of an operation such as a *command*. It can be used to synchronize operations in a context.

Event Wait List: An *event wait list* is a list of *event objects* that can be used to control when a particular *command* begins execution.

Framework: A software system that contains the set of components to support software development and execution. A *framework* typically includes libraries, APIs, runtime systems, compilers, etc.

Global ID: 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*. The *global ID* is a N-dimensional value that starts at (0, 0, ... 0). See also *Local ID*.

Global Memory: A memory region accessible to all *work-items* executing in a *context*. It is accessible to the *host* using *commands* such as read, write and map.

GL share group: A *GL share group* object manages shared OpenGL or OpenGL ES resources such as textures, buffers, framebuffers, and renderbuffers and is associated with one or more GL context objects. The *GL share group* is typically an opaque object and not directly accessible.

Handle: An opaque type that references an *object* allocated by OpenCL. Any operation on an *object* occurs by reference to that object's handle.

Host: The *host* interacts with the *context* using the OpenCL API.

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

Illegal: Behavior of a system that is explicitly not allowed and will be reported as an error when encountered by OpenCL.

Image Object: A *memory object* that stores a two- or three- dimensional structured array. Image data can only be accessed with read and write functions. The read functions use a *sampler*.

The *image object* encapsulates the following information:

- ✚ Dimensions of the image.
- ✚ Description of each element in the image.
- ✚ Properties that describe usage information and which region to allocate from.
- ✚ Image data.

The elements of an image are selected from a list of predefined image formats.

Implementation Defined: Behavior that is explicitly allowed to vary between conforming implementations of OpenCL. An OpenCL implementor is required to document the implementation-defined behavior.

In-order Execution: A model of execution in OpenCL where the *commands* in a *command-queue* are executed in order of submission with each *command* running to completion before the next one begins. See *Out-of-order Execution*.

Kernel: A *kernel* is a function declared in a *program* and executed on an OpenCL *device*. A *kernel* is identified by the `__kernel` or `kernel` qualifier applied to any function defined in a *program*.

Kernel Object: A *kernel object* encapsulates a specific `__kernel` function declared in a *program* and the argument values to be used when executing this `__kernel` function.

Local ID: A *local ID* specifies a unique *work-item ID* within a given *work-group* that is executing a *kernel*. The *local ID* is a N-dimensional value that starts at (0, 0, ... 0). See also *Global ID*.

Local Memory: A memory region associated with a *work-group* and accessible only by *work-items* in that *work-group*.

Marker: A *command* queued in a *command-queue* that can be used to tag all *commands* queued before the *marker* in the *command-queue*. The *marker* command returns an *event* which can be used by the *application* to queue a wait on the marker event i.e. wait for all commands queued before the *marker* command to complete.

Memory Objects: A *memory object* is a handle to a reference counted region of *global memory*. Also see *Buffer Object* and *Image Object*.

Memory Regions (or Pools): A distinct address space in OpenCL. *Memory regions* may overlap in physical memory though OpenCL will treat them as logically distinct. The *memory regions* are denoted as *private*, *local*, *constant*, and *global*.

Object: Objects are abstract representation of the resources that can be manipulated by the OpenCL API. Examples include *program objects*, *kernel objects*, and *memory objects*.

Out-of-Order Execution: A model of execution in which *commands* placed in the *work queue* may begin and complete execution in any order consistent with constraints imposed by *event wait lists* and *command-queue barrier*. See *In-order Execution*.

Parent device: The OpenCL *device* which is partitioned to create *sub-devices*. Not all *parent devices* are *root devices*. A *root device* might be partitioned and the *sub-devices* partitioned again. In this case, the first set of *sub-devices* would be *parent devices* of the second set, but not the *root devices*. Also see *device*, *parent device* and *root device*.

Platform: The *host* plus a collection of *devices* managed by the OpenCL *framework* that allow an application to share *resources* and execute *kernels* on *devices* in the *platform*.

Private Memory: A region of memory private to a *work-item*. Variables defined in one *work-item*'s private memory are not visible to another *work-item*.

Processing Element: A virtual scalar processor. A work-item may execute on one or more processing elements.

Program: An OpenCL *program* consists of a set of *kernels*. *Programs* may also contain auxiliary functions called by the `__kernel` functions and constant data.

Program Object: A *program object* encapsulates the following information:

- ✚ A reference to an associated *context*.
- ✚ A *program* source or binary.
- ✚ The latest successfully built program executable, the list of *devices* for which the program executable is built, the build options used and a build log.
- ✚ The number of *kernel objects* currently attached.

Reference Count: The life span of an OpenCL object is determined by its *reference count*—an internal count of the number of references to the object. When you create an object in OpenCL, its *reference count* is set to one. Subsequent calls to the appropriate *retain* API (such as `clRetainContext`, `clRetainCommandQueue`) increment the *reference count*. Calls to the appropriate *release* API (such as `clReleaseContext`, `clReleaseCommandQueue`) decrement the *reference count*. After the *reference count* reaches zero, the object's resources are deallocated by OpenCL.

Relaxed Consistency: A memory consistency model in which the contents of memory visible to different *work-items* or *commands* may be different except at a *barrier* or other explicit synchronization points.

Resource: A class of *objects* defined by OpenCL. An instance of a *resource* is an *object*. The most common *resources* are the *context*, *command-queue*, *program objects*, *kernel objects*, and *memory objects*. Computational resources are hardware elements that participate in the action of advancing a program counter. Examples include the *host*, *devices*, *compute units* and *processing elements*.

Retain, Release: The action of incrementing (retain) and decrementing (release) the reference count using an OpenCL *object*. This is a book keeping functionality to make sure the system doesn't remove an *object* before all instances that use this *object* have finished. Refer to *Reference Count*.

Root device: A *root device* is an OpenCL *device* that has not been partitioned. Also see *device*, *parent device* and *root device*.

Sampler: An *object* that describes how to sample an image when the image is read in the *kernel*. The image read functions take a *sampler* as an argument. The *sampler* specifies the image

addressing-mode i.e. how out-of-range image coordinates are handled, the filter mode, and whether the input image coordinate is a normalized or unnormalized value.

SIMD: Single Instruction Multiple Data. A programming model where a *kernel* is executed concurrently on multiple *processing elements* each with its own data and a shared program counter. All *processing elements* execute a strictly identical set of instructions.

SPMD: Single Program Multiple Data. A programming model where a *kernel* is executed concurrently on multiple *processing elements* each with its own data and its own program counter. Hence, while all computational resources run the same *kernel* they maintain their own instruction counter and due to branches in a *kernel*, the actual sequence of instructions can be quite different across the set of *processing elements*.

Sub-device: An OpenCL *device* can be partitioned into multiple *sub-devices*. The new *sub-devices* alias specific collections of compute units within the parent *device*, according to a partition scheme. The *sub-devices* may be used in any situation that their parent *device* may be used. Partitioning a *device* does not destroy the parent *device*, which may continue to be used along side and intermingled with its child *sub-devices*. Also see *device*, *parent device* and *root device*.

Task Parallel Programming Model: A programming model in which computations are expressed in terms of multiple concurrent tasks where a task is a *kernel* executing in a single *work-group* of size one. The concurrent tasks can be running different *kernels*.

Thread-safe: An OpenCL API call is considered to be *thread-safe* if the internal state as managed by OpenCL remains consistent when called simultaneously by multiple *host* threads. OpenCL API calls that are *thread-safe* allow an application to call these functions in multiple *host* threads without having to implement mutual exclusion across these *host* threads i.e. they are also re-entrant-safe.

Undefined: The behavior of an OpenCL API call, built-in function used inside a *kernel* or execution of a *kernel* that is explicitly not defined by OpenCL. A conforming implementation is not required to specify what occurs when an undefined construct is encountered in OpenCL.

Work-group: 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-group Barrier. See *Barrier*.

Work-item: 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*.

2.1 OpenCL Class Diagram

Figure 2.1 describes the OpenCL specification as a class diagram using the Unified Modeling Language¹ (UML) notation. The diagram shows both nodes and edges which are classes and their relationships. As a simplification it shows only classes, and no attributes or operations. Abstract classes are annotated with “{abstract}”. As for relationships it shows aggregations (annotated with a solid diamond), associations (no annotation), and inheritance (annotated with an open arrowhead). The cardinality of a relationship is shown on each end of it. A cardinality of “*” represents “many”, a cardinality of “1” represents “one and only one”, a cardinality of “0..1” represents “optionally one”, and a cardinality of “1..*” represents “one or more”. The navigability of a relationship is shown using a regular arrowhead.

¹ Unified Modeling Language (<http://www.uml.org/>) is a trademark of Object Management Group (OMG).

3. The OpenCL Architecture

OpenCL is an open industry standard for programming a heterogeneous collection of CPUs, GPUs and other discrete computing devices organized into a single platform. It is more than a language. OpenCL is a framework for parallel programming and includes a language, API, libraries and a runtime system to support software development. Using OpenCL, for example, a programmer can write general purpose programs that execute on GPUs without the need to map their algorithms onto a 3D graphics API such as OpenGL or DirectX.

The target of OpenCL is expert programmers wanting to write portable yet efficient code. This includes library writers, middleware vendors, and performance oriented application programmers. Therefore OpenCL provides a low-level hardware abstraction plus a framework to support programming and many details of the underlying hardware are exposed.

To describe the core ideas behind OpenCL, we will use a hierarchy of models:

- ✚ Platform Model
- ✚ Memory Model
- ✚ Execution Model
- ✚ Programming Model

3.1 Platform Model

The Platform model for OpenCL is defined in *figure 3.1*. 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 further divided into one or more **processing elements** (PEs). Computations on a device occur within the processing elements.

An OpenCL application runs on a host according to the models native to the host platform. The OpenCL application submits **commands** from the host to execute computations on the processing elements within a device. The processing elements within a compute unit execute a single stream of instructions as SIMD units (execute in lockstep with a single stream of instructions) or as SPMD units (each PE maintains its own program counter).




Figure 3.1: *Platform model ... one host plus one or more compute devices each with one or more compute units each with one or more processing elements.*

3.1.1 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. There are three important version identifiers to consider for an OpenCL system: the platform version, the version of a device, and the version(s) of the OpenCL C language supported on a device.

The platform version indicates the version of the OpenCL runtime supported. This includes all of the APIs that the host can use to interact with the OpenCL runtime, such as contexts, memory objects, devices, and command queues.

The device version is an indication of the devices capabilities, separate from the runtime and compiler, as represented by the device info returned by `clGetDeviceInfo`. Examples of attributes associated with the device version are resource limits and extended functionality. The version returned corresponds to the highest version of the OpenCL spec for which the device is conformant, but is not higher than the platform version.

The language version for a device represents the OpenCL programming language features a developer can assume are supported on a given device. The version reported is the highest version of the language supported.

OpenCL C is designed to be backwards compatible, so a device is not required to support more than a single language version to be considered conformant. If multiple language versions are supported, the compiler defaults to using the highest language version supported for the device. The language version is not higher than the platform version, but may exceed the device version (see *section 5.6.4.5*).

3.2 Execution Model

Execution of an OpenCL program occurs in two parts: **kernels** that execute on one or more **OpenCL devices** and a **host program** that executes on the host. The host program defines the context for the kernels and manages their execution.

The core of the OpenCL execution model is defined by how the kernels execute. When a kernel is submitted for execution by the host, an index space is defined. An instance of the kernel 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 per work-item.

Work-items are organized into **work-groups**. The work-groups provide a more coarse-grained decomposition of the index space. Work-groups are assigned a unique work-group ID with the same dimensionality as the index space used for the work-items. Work-items are assigned a unique local ID within a 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 OpenCL is called an NDRange. An NDRange is an N-dimensional index space, where N is one, two or three. An NDRange is defined by an integer array of length N specifying the extent of the index space in each dimension starting at an offset index F (zero by default). Each work-item's global ID and local ID are N-dimensional tuples. The global ID components are values in the range from F, to F plus the number of elements in that dimension minus one.

Work-groups are assigned IDs using a similar approach to that used for work-item global IDs. An array of length N defines the number of work-groups in each dimension. 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. Each work-item is identifiable in two ways; in terms of a global index, and in terms of a work-group index plus a local index within a work group.

For example, consider the 2-dimensional index space in *figure 3.2*. We input the index space for the work-items (G_x, G_y), the size of each work-group (S_x, S_y) and the global ID offset (F_x, F_y). The global indices define an G_x by G_y index space where the total number of work-items is the product of G_x and G_y . The local indices define a S_x by S_y index space where the number of work-items in a single work-group is the product of S_x and S_y . Given the size of each work-group and the total number of work-items we can compute the number of work-groups. A 2-dimensional index space is used to uniquely identify a work-group. Each work-item is identified by its global ID (g_x, g_y) or by the combination of the work-group ID (w_x, w_y), the size of each work-group (S_x, S_y) and the local ID (s_x, s_y) inside the work-group such that

$$(g_x, g_y) = (w_x * S_x + s_x + F_x, w_y * S_y + s_y + F_y)$$

The number of work-groups can be computed as:

$$(W_x, W_y) = (G_x / S_x, G_y / S_y)$$

Given a global ID and the work-group size, the work-group ID for a work-item is computed as:

$$(w_x, w_y) = ((g_x - s_x - F_x) / S_x, (g_y - s_y - F_y) / S_y)$$

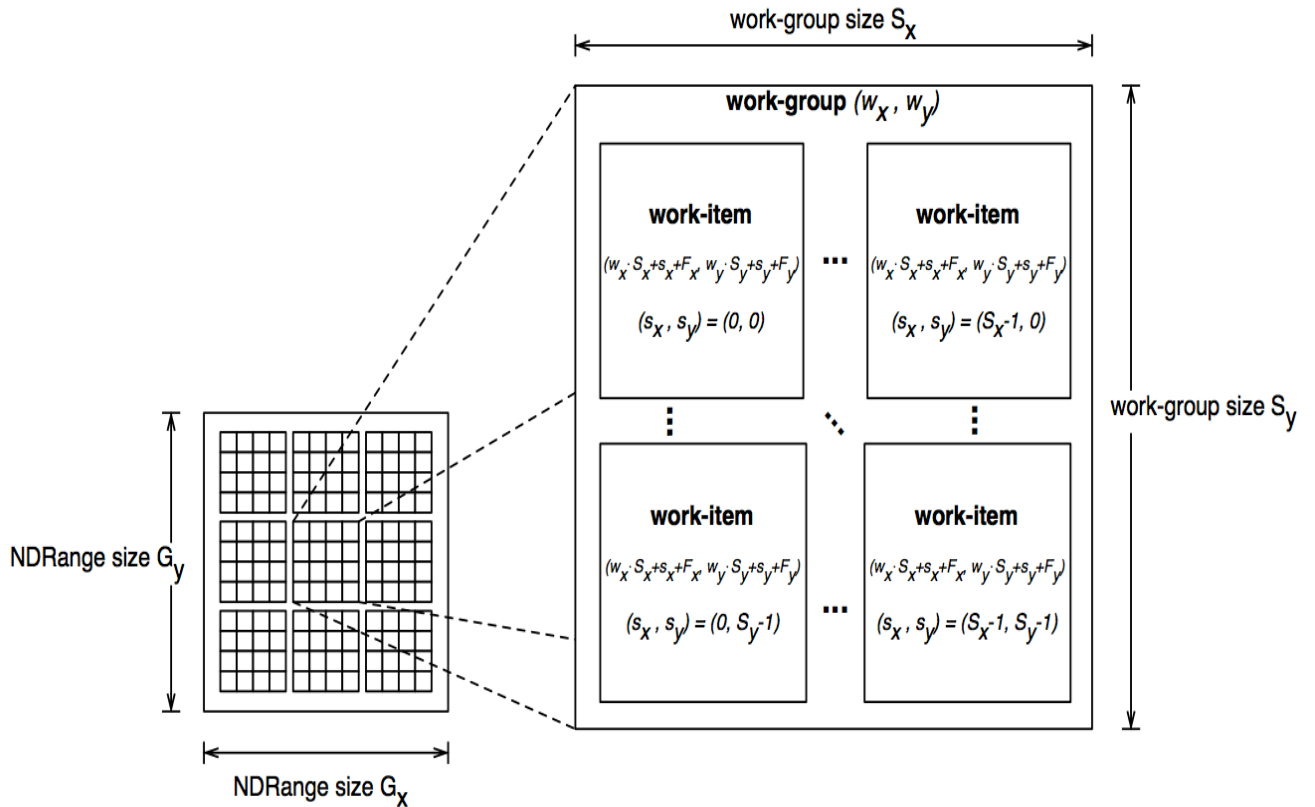


Figure 3.2 An example of an NDRange index space showing work-items, their global IDs and their mapping onto the pair of work-group and local IDs.

A wide range of programming models can be mapped onto this execution model. We explicitly support two of these models within OpenCL; the **data parallel programming model** and the **task parallel programming model**.

3.2.1 Execution Model: Context and Command Queues

The host defines a context for the execution of the kernels. The context includes the following resources:

1. **Devices:** The collection of OpenCL devices to be used by the host.
2. **Kernels:** The OpenCL functions that run on OpenCL devices.
3. **Program Objects:** The program source and executable that implement the kernels.
4. **Memory Objects:** A set of memory objects visible to the host and the OpenCL devices. Memory objects contain values that can be operated on by instances of a kernel.

The context is created and manipulated by the host using functions from the OpenCL API. The host creates a data structure called a **command-queue** to coordinate execution of the kernels on the devices. The host places commands into the command-queue which are then scheduled onto the devices within the context. These include:

- ✚ **Kernel execution commands:** Execute a kernel on the processing elements of a device.
- ✚ **Memory commands:** Transfer data to, from, or between memory objects, or map and unmap memory objects from the host address space.
- ✚ **Synchronization commands:** Constrain the order of execution of commands.

The command-queue schedules commands for execution on a device. These execute asynchronously between the host and the device. Commands execute relative to each other in one of two modes:

- ✚ **In-order Execution:** Commands are launched in the order they appear in the command-queue and complete in order. In other words, a prior command on the queue completes before the following command begins. This serializes the execution order of commands in a queue.
- ✚ **Out-of-order Execution:** Commands are issued in order, but do not wait to complete before following commands execute. Any order constraints are enforced by the programmer through explicit synchronization commands.

Kernel execution and memory commands submitted to a queue generate event objects. These are used to control execution between commands and to coordinate execution between the host and devices.

It is possible to associate multiple queues with a single context. These queues run concurrently and independently with no explicit mechanisms within OpenCL to synchronize between them.

3.2.2 Execution Model: Categories of Kernels

The OpenCL execution model supports two categories of kernels:

- ✚ **OpenCL kernels** are written with the OpenCL C programming language and compiled with the OpenCL compiler. All OpenCL implementations support OpenCL kernels. Implementations may provide other mechanisms for creating OpenCL kernels.
- ✚ **Native kernels** are accessed through a host function pointer. Native kernels are queued for execution along with OpenCL kernels on a device and share memory objects with OpenCL kernels. For example, these native kernels could be functions defined in application code or exported from a library. Note that the ability to execute native kernels is an optional functionality within OpenCL and the semantics of native kernels are implementation-defined. The OpenCL API includes functions to query capabilities of a device(s) and determine if this capability is supported.

3.3 Memory Model

Work-item(s) executing a kernel have access to four distinct memory regions:

- ✚ **Global Memory.** This memory region permits read/write access to all work-items in all work-groups. Work-items can read from or write to any element of a memory object. Reads and writes to global memory may be cached depending on the capabilities of the device.
- ✚ **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.
- ✚ **Local Memory:** A memory region local to a work-group. This memory region can be used to allocate variables that are shared by all work-items in that work-group. It may be implemented as dedicated regions of memory on the OpenCL device. Alternatively, the local memory region may be mapped onto sections of the global memory.
- ✚ **Private Memory:** A region of memory private to a work-item. Variables defined in one work-item's private memory are not visible to another work-item.

Table 3.1 describes whether the kernel or the host can allocate from a memory region, the type of allocation (static i.e. compile time vs dynamic i.e. runtime) and the type of access allowed i.e. whether the kernel or the host can read and/or write to a memory region.

| | Global | Constant | Local | Private |
|---------------|---------------------|---------------------|---------------------|---------------------|
| Host | Dynamic allocation | Dynamic allocation | Dynamic allocation | No allocation |
| | Read / Write access | Read / Write access | No access | No access |
| Kernel | No allocation | Static allocation | Static allocation | Static allocation |
| | Read / Write access | Read-only access | Read / Write access | Read / Write access |

Table 3.1 *Memory Region - Allocation and Memory Access Capabilities*

The memory regions and how they relate to the platform model are described in *figure 3.3*.

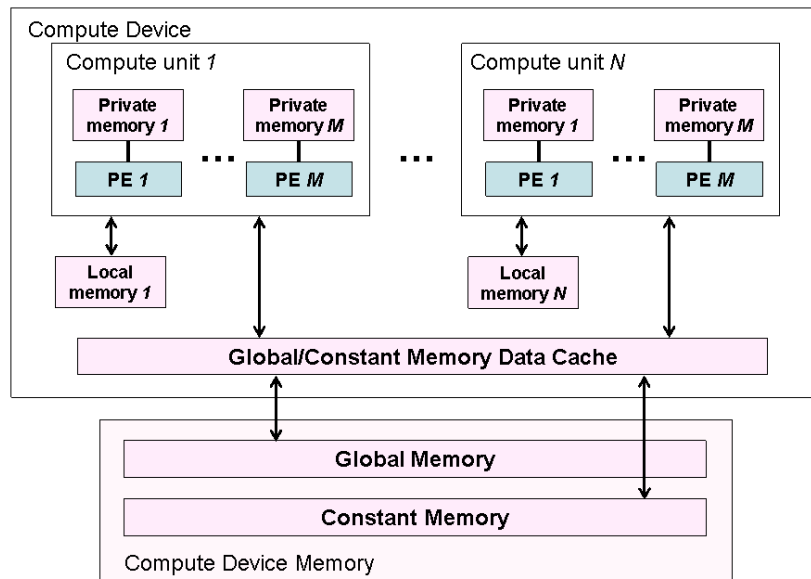


Figure 3.3: *Conceptual OpenCL device architecture with processing elements (PE), compute units and devices. The host is not shown.*

The application running on the host uses the OpenCL API to create memory objects in global memory, and to enqueue memory commands (described in *section 3.2.1*) that operate on these memory objects.

The host and OpenCL device memory models are, for the most part, independent of each other. This is by a necessity given that the host is defined outside of OpenCL. They do, however, at times need to interact. This interaction occurs in one of two ways: by explicitly copying data or by mapping and unmapping regions of a memory object.

To copy data explicitly, the host enqueues commands to transfer data between the memory object and host memory. These memory transfer commands may be blocking or non-blocking. The OpenCL function call for a blocking memory transfer returns once the associated memory resources on the host can be safely reused. For a non-blocking memory transfer, the OpenCL function call returns as soon as the command is enqueued regardless of whether host memory is safe to use.

The mapping/unmapping method of interaction between the host and OpenCL memory objects allows the host to map a region from the memory object into its address space. The memory map command may be blocking or non-blocking. Once a region from the memory object has been mapped, the host can read or write to this region. The host unmaps the region when accesses (reads and/or writes) to this mapped region by the host are complete.

3.3.1 Memory Consistency

OpenCL uses a relaxed consistency memory 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.

Within a work-item memory has load / store consistency. Local memory is consistent across work-items in a single work-group at a work-group barrier. Global memory is consistent across work-items in a single work-group at a work-group barrier, but there are no guarantees of memory consistency between different work-groups executing a kernel.

Memory consistency for memory objects shared between enqueued commands is enforced at a synchronization point.

3.4 Programming Model

The OpenCL execution model supports **data parallel** and **task parallel** programming models, as well as supporting hybrids of these two models. The primary model driving the design of OpenCL is data parallel.

3.4.1 Data Parallel Programming Model

A data parallel programming model defines a computation in terms of a sequence of instructions applied to multiple elements of a memory object. The index space associated with the OpenCL execution model defines the work-items and how the data maps onto the work-items. In a strictly data parallel model, there is a one-to-one mapping between the work-item and the element in a memory object over which a kernel can be executed in parallel. OpenCL implements a relaxed version of the data parallel programming model where a strict one-to-one mapping is not a requirement.

OpenCL provides a hierarchical data parallel programming model. There are two ways to specify the hierarchical subdivision. In the explicit model a programmer defines the total number of work-items to execute in parallel and also how the work-items are divided among work-groups. In the implicit model, a programmer specifies only the total number of work-items to execute in parallel, and the division into work-groups is managed by the OpenCL implementation.

3.4.2 Task Parallel Programming Model

The OpenCL task parallel programming model defines a model in which a single instance of a kernel is executed independent of any index space. It is logically equivalent to executing a kernel on a compute unit with a work-group containing a single work-item. Under this model, users express parallelism by:

- using vector data types implemented by the device,
- enqueueing multiple tasks, and/or
- enqueueing native kernels developed using a programming model orthogonal to OpenCL.

3.4.3 Synchronization

There are two domains of synchronization in OpenCL:

- Work-items in a single work-group
- Commands enqueued to command-queue(s) in a single context

Synchronization between work-items in a single work-group is done using a work-group barrier. 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.

The synchronization points between commands in command-queues are:

- Command-queue barrier. The command-queue barrier ensures that all previously queued commands have finished execution and any resulting updates to memory objects are visible to subsequently enqueued commands before they begin execution. This barrier can only be used to synchronize between commands in a single command-queue.
- Waiting on an event. All OpenCL API functions that enqueue commands return an event that identifies the command and memory objects it updates. A subsequent command waiting on that event is guaranteed that updates to those memory objects are visible before the command begins execution.

3.5 Memory Objects

Memory objects are categorized into two types: *buffer* objects, and *image* objects. A *buffer* object stores a one-dimensional collection of elements whereas an *image* object is used to store a two- or three- dimensional texture, frame-buffer or image.

Elements of a *buffer* object can be a scalar data type (such as an int, float), vector data type, or a user-defined structure. An *image* object is used to represent a buffer that can be used as a texture or a frame-buffer. The elements of an image object are selected from a list of predefined image formats. 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 sequential fashion and can be accessed using a pointer by a kernel executing on a device. Elements of an *image* are stored in a format that is opaque to the user and cannot be directly accessed using a pointer. Built-in functions are provided by the OpenCL C programming language 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 built-in function to read from an image converts image element from the format it is stored into a 4-component vector. Similarly, the built-in function to write to an image converts the image element from a 4-component vector to the appropriate image format specified such as 4 8-bit elements, for example.

Memory objects are described by a **cl_mem** object. Kernels take memory objects as input, and output to one or more memory objects.

3.6 The OpenCL Framework

The OpenCL 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:

- ✚ **OpenCL Platform layer:** The platform layer allows the host program to discover OpenCL devices and their capabilities and to create contexts.

- ✚ **OpenCL Runtime:** The runtime allows the host program to manipulate contexts once they have been created.
- ✚ **OpenCL Compiler:** The OpenCL compiler creates program executables that contain OpenCL kernels. The OpenCL C programming language implemented by the compiler supports a subset of the ISO C99 language with extensions for parallelism.

4. The OpenCL Platform Layer

This section describes the OpenCL platform layer which implements platform-specific features that allow applications to query OpenCL devices, device configuration information, and to create OpenCL contexts using one or more devices.

4.1 Querying Platform Info

The list of platforms available can be obtained using the following function.

```
cl_int  clGetPlatformIDs (cl_uint num_entries,  
                        cl_platform_id *platforms,  
                        cl_uint *num_platforms)
```

num_entries is the number of *cl_platform_id* entries that can be added to *platforms*. If *platforms* is not NULL, the *num_entries* must be greater than zero.

platforms returns a list of OpenCL platforms found. The *cl_platform_id* values returned in *platforms* can be used to identify a specific OpenCL platform. If *platforms* argument is NULL, this argument is ignored. The number of OpenCL platforms returned is the minimum of the value specified by *num_entries* or the number of OpenCL platforms available.

num_platforms returns the number of OpenCL platforms available. If *num_platforms* is NULL, this argument is ignored.

clGetPlatformIDs returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_VALUE if *num_entries* is equal to zero and *platforms* is not NULL or if both *num_platforms* and *platforms* are NULL.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```
cl_int  clGetPlatformInfo (cl_platform_id platform,  
                        cl_platform_info param_name,  
                        size_t param_value_size,  
                        void *param_value,  
                        size_t *param_value_size_ret)
```

gets specific information about the OpenCL platform. The information that can be queried using **clGetPlatformInfo** is specified in *table 4.1*.

platform refers to the platform ID returned by **clGetPlatformIDs** or can be NULL. If *platform* is NULL, the behavior is implementation-defined.

param_name is an enumeration constant that identifies the platform information being queried. It can be one of the following values as specified in *table 4.1*.

param_value is a pointer to memory location where appropriate values for a given *param_name* as specified in *table 4.1* will be returned. If *param_value* is NULL, it is ignored.

param_value_size specifies the size in bytes of memory pointed to by *param_value*. This size in bytes must be \geq size of return type specified in *table 4.1*.

param_value_size_ret returns the actual size in bytes of data being queried by *param_value*. If *param_value_size_ret* is NULL, it is ignored.

| cl_platform_info | Return Type | Description |
|----------------------------|---------------------|---|
| CL_PLATFORM_PROFILE | char[] ² | OpenCL profile string. Returns the profile name supported by the implementation. The profile name returned can be one of the following strings: FULL_PROFILE – if the implementation supports the OpenCL specification (functionality defined as part of the core specification and does not require any extensions to be supported). EMBEDDED_PROFILE - if the implementation supports the OpenCL embedded profile. The embedded profile is defined to be a subset for each version of OpenCL. The embedded profile for OpenCL 1.2 is described in <i>section 10</i> . |
| CL_PLATFORM_VERSION | char[] | OpenCL version string. Returns the OpenCL version supported by the implementation. This version string has the following format: |

² A null terminated string is returned by OpenCL query function calls if the return type of the information being queried is a char[].

| | | |
|-------------------------------|--------|---|
| | | <i>OpenCL</i> <space><major_version.minor_version><space><platform-specific information> The <i>major_version.minor_version</i> value returned will be 1.2. |
| CL_PLATFORM_NAME | char[] | Platform name string. |
| CL_PLATFORM_VENDOR | char[] | Platform vendor string. |
| CL_PLATFORM_EXTENSIONS | char[] | Returns a space separated list of extension names (the extension names themselves do not contain any spaces) supported by the platform. Extensions defined here must be supported by all devices associated with this platform. |

Table 4.1. *OpenCL Platform Queries*

clGetPlatformInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors³:

- ✚ CL_INVALID_PLATFORM if *platform* is not a valid platform.
- ✚ CL_INVALID_VALUE if *param_name* is not one of the supported values or if size in bytes specified by *param_value_size* is < size of return type as specified in *table 4.1* and *param_value* is not a NULL value.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

4.2 Querying Devices

The list of devices available on a platform can be obtained using the following function.

```
cl_int          clGetDeviceIDs4(cl_platform_id platform,
                                cl_device_type device_type,
                                cl_uint num_entries,
                                cl_device_id *devices,
                                cl_uint *num_devices)
```

³ The OpenCL specification does not describe the order of precedence for error codes returned by API calls.

⁴ **clGetDeviceIDs** may return all or a subset of the actual physical devices present in the platform and that match *device_type*.

platform refers to the platform ID returned by **clGetPlatformIDs** or can be NULL. If *platform* is NULL, the behavior is implementation-defined.

device_type is a bitfield that identifies the type of OpenCL device. The *device_type* can be used to query specific OpenCL devices or all OpenCL devices available. The valid values for *device_type* are specified in *table 4.2*.

| cl_device_type | Description |
|-----------------------------------|---|
| CL_DEVICE_TYPE_CPU | An OpenCL device that is the host processor. The host processor runs the OpenCL implementations and is a single or multi-core CPU. |
| CL_DEVICE_TYPE_GPU | An OpenCL device that is a GPU. By this we mean that the device can also be used to accelerate a 3D API such as OpenGL or DirectX. |
| CL_DEVICE_TYPE_ACCELERATOR | Dedicated OpenCL accelerators (for example the IBM CELL Blade). These devices communicate with the host processor using a peripheral interconnect such as PCIe. |
| CL_DEVICE_TYPE_CUSTOM | Dedicated accelerators that do not support programs written in OpenCL C. |
| CL_DEVICE_TYPE_DEFAULT | The default OpenCL device in the system. The default device cannot be a CL_DEVICE_TYPE_CUSTOM device. |
| CL_DEVICE_TYPE_ALL | All OpenCL devices available in the system except CL_DEVICE_TYPE_CUSTOM devices.. |


Table 4.2. *List of OpenCL Device Categories*

num_entries is the number of *cl_device_id* entries that can be added to *devices*. If *devices* is not NULL, the *num_entries* must be greater than zero.

devices returns a list of OpenCL devices found. The *cl_device_id* values returned in *devices* can be used to identify a specific OpenCL device. If *devices* argument is NULL, this argument is ignored. The number of OpenCL devices returned is the minimum of the value specified by *num_entries* or the number of OpenCL devices whose type matches *device_type*.

num_devices returns the number of OpenCL devices available that match *device_type*. If *num_devices* is NULL, this argument is ignored.

clGetDeviceIDs returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

-  CL_INVALID_PLATFORM if *platform* is not a valid platform.

- ✚ CL_INVALID_DEVICE_TYPE if *device_type* is not a valid value.
- ✚ CL_INVALID_VALUE if *num_entries* is equal to zero and *devices* is not NULL or if both *num_devices* and *devices* are NULL.
- ✚ CL_DEVICE_NOT_FOUND if no OpenCL devices that matched *device_type* were found.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The application can query specific capabilities of the OpenCL device(s) returned by **clGetDeviceIDs**. This can be used by the application to determine which device(s) to use.

The function

```

cl_int          clGetDeviceInfo (cl_device_id device,
                                   cl_device_info param_name,
                                   size_t param_value_size,
                                   void *param_value,
                                   size_t *param_value_size_ret)

```

gets specific information about an OpenCL device. *device* may be a device returned by **clGetDeviceIDs** or a sub-device created by **clCreateSubDevices**. If *device* is a sub-device, the specific information for the sub-device will be returned. The information that can be queried using **clGetDeviceInfo** is specified in *table 4.3*.

device is a device returned by **clGetDeviceIDs**.

param_name is an enumeration constant that identifies the device information being queried. It can be one of the following values as specified in *table 4.3*.

param_value is a pointer to memory location where appropriate values for a given *param_name* as specified in *table 4.3* will be returned. If *param_value* is NULL, it is ignored.

param_value_size specifies the size in bytes of memory pointed to by *param_value*. This size in bytes must be \geq size of return type specified in *table 4.3*.

param_value_size_ret returns the actual size in bytes of data being queried by *param_value*. If *param_value_size_ret* is NULL, it is ignored.

| cl_device_info | Return Type | Description |
|---|--------------------|---|
| CL_DEVICE_TYPE | cl_device_type | The OpenCL device type. Currently supported values are: CL_DEVICE_TYPE_CPU, CL_DEVICE_TYPE_GPU, CL_DEVICE_TYPE_ACCELERATOR, CL_DEVICE_TYPE_DEFAULT, a combination of the above types or CL_DEVICE_TYPE_CUSTOM. |
| CL_DEVICE_VENDOR_ID | cl_uint | A unique device vendor identifier. An example of a unique device identifier could be the PCIe ID. |
| CL_DEVICE_MAX_COMPUTE_UNITS | cl_uint | The number of parallel compute units on the OpenCL device. A work-group executes on a single compute unit. The minimum value is 1. |
| CL_DEVICE_MAX_WORK_ITEM_DIMENSIONS | cl_uint | Maximum dimensions that specify the global and local work-item IDs used by the data parallel execution model. (Refer to clEnqueueNDRangeKernel). The minimum value is 3 for devices that are not of type CL_DEVICE_TYPE_CUSTOM. |
| CL_DEVICE_MAX_WORK_ITEM_SIZES | size_t [] | Maximum number of work-items that can be specified in each dimension of the work-group to clEnqueueNDRangeKernel . Returns <i>n</i> size_t entries, where <i>n</i> is the value returned by the query for CL_DEVICE_MAX_WORK_ITEM_DIMENSIONS. The minimum value is (1, 1, 1) for devices that are not of type CL_DEVICE_TYPE_CUSTOM. |
| CL_DEVICE_MAX_WORK_GROUP_SIZE | size_t | Maximum number of work-items in a work-group executing a kernel on a single compute unit, using the data parallel execution model. (Refer to clEnqueueNDRangeKernel). The minimum value is 1. |
| CL_DEVICE_PREFERRED_VECTOR_WIDTH_CHAR CL_DEVICE_PREFERRED_VECTOR_WIDTH_SHORT | cl_uint | Preferred native vector width size for built-in 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 |

| | | |
|---|----------|--|
| <p>CL_DEVICE_PREFERRED_VECTOR_WIDTH_INT</p> <p>CL_DEVICE_PREFERRED_VECTOR_WIDTH_LONG</p> <p>CL_DEVICE_PREFERRED_VECTOR_WIDTH_FLOAT</p> <p>CL_DEVICE_PREFERRED_VECTOR_WIDTH_DOUBLE</p> <p>CL_DEVICE_PREFERRED_VECTOR_WIDTH_HALF</p> | | <p>vector.</p> <p>If double precision is not supported, CL_DEVICE_PREFERRED_VECTOR_WIDTH_DOUBLE must return 0.</p> <p>If the cl_khr_fp16 extension is not supported, CL_DEVICE_PREFERRED_VECTOR_WIDTH_HALF must return 0.</p> |
| <p>CL_DEVICE_NATIVE_VECTOR_WIDTH_CHAR</p> <p>CL_DEVICE_NATIVE_VECTOR_WIDTH_SHORT</p> <p>CL_DEVICE_NATIVE_VECTOR_WIDTH_INT</p> <p>CL_DEVICE_NATIVE_VECTOR_WIDTH_LONG</p> <p>CL_DEVICE_NATIVE_VECTOR_WIDTH_FLOAT</p> <p>CL_DEVICE_NATIVE_VECTOR_WIDTH_DOUBLE</p> <p>CL_DEVICE_NATIVE_VECTOR_WIDTH_HALF</p> | cl_uint | <p>Returns the native ISA vector width. The vector width is defined as the number of scalar elements that can be stored in the vector.</p> <p>If double precision is not supported, CL_DEVICE_NATIVE_VECTOR_WIDTH_DOUBLE must return 0.</p> <p>If the cl_khr_fp16 extension is not supported, CL_DEVICE_NATIVE_VECTOR_WIDTH_HALF must return 0.</p> |
| CL_DEVICE_MAX_CLOCK_FREQUENCY | cl_uint | Maximum configured clock frequency of the device in MHz. |
| CL_DEVICE_ADDRESS_BITS | cl_uint | The default compute device address space size specified as an unsigned integer value in bits. Currently supported values are 32 or 64 bits. |
| CL_DEVICE_MAX_MEM_ALLOC_SIZE | cl_ulong | Max size of memory object allocation in bytes. The minimum value is max (1/4 th of CL_DEVICE_GLOBAL_MEM_SIZE , 128*1024*1024) for devices that are not of type CL_DEVICE_TYPE_CUSTOM . |
| CL_DEVICE_IMAGE_SUPPORT | cl_bool | Is CL_TRUE if images are supported by the OpenCL device and CL_FALSE otherwise. |
| CL_DEVICE_MAX_READ_IMAGE_ARGS | cl_uint | Max number of simultaneous image objects that can be read by a kernel. The minimum value is 128 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE . |

| | | |
|--|---------|---|
| CL_DEVICE_MAX_WRITE_IMAGE_ARGS | cl_uint | Max number of simultaneous image objects that can be written to by a kernel. The minimum value is 8 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_IMAGE2D_MAX_WIDTH | size_t | Max width of 2D image or 1D image not created from a buffer object in pixels. The minimum value is 8192 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_IMAGE2D_MAX_HEIGHT | size_t | Max height of 2D image in pixels. The minimum value is 8192 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_IMAGE3D_MAX_WIDTH | size_t | Max width of 3D image in pixels. The minimum value is 2048 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_IMAGE3D_MAX_HEIGHT | size_t | Max height of 3D image in pixels. The minimum value is 2048 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_IMAGE3D_MAX_DEPTH | size_t | Max depth of 3D image in pixels. The minimum value is 2048 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_IMAGE_MAX_BUFFER_SIZE | size_t | Max number of pixels for a 1D image created from a buffer object. The minimum value is 65536 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_IMAGE_MAX_ARRAY_SIZE | size_t | Max number of images in a 1D or 2D image array. The minimum value is 2048 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_MAX_SAMPLERS | cl_uint | Maximum number of samplers that can be used in a kernel. Refer to <i>section 6.12.14</i> for a detailed description on samplers. |

| | | |
|--------------------------------------|---------------------|---|
| | | The minimum value is 16 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_MAX_PARAMETER_SIZE | size_t | 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 CL_DEVICE_TYPE_CUSTOM. For this minimum value, only a maximum of 128 arguments can be passed to a kernel. |
| CL_DEVICE_MEM_BASE_ADDR_ALIGN | cl_uint | The minimum value is the size (in bits) of the largest OpenCL built-in data type supported by the device (long16 in FULL profile, long16 or int16 in EMBEDDED profile) for devices that are not of type CL_DEVICE_TYPE_CUSTOM. |
| CL_DEVICE_SINGLE_FP_CONFIG | cl_device_fp_config | Describes single precision floating-point capability of the device. This is a bit-field that describes one or more of the following values: CL_FP_DENORM – denorms are supported CL_FP_INF_NAN – INF and quiet NaNs are supported. CL_FP_ROUND_TO_NEAREST – round to nearest even rounding mode supported CL_FP_ROUND_TO_ZERO – round to zero rounding mode supported CL_FP_ROUND_TO_INF – round to positive and negative infinity rounding modes supported CL_FP_FMA – IEEE754-2008 fused multiply-add is supported. CL_FP_CORRECTLY_ROUNDED_DIVIDE_SQRT – divide and sqrt are correctly rounded as defined by the IEEE754 specification. |

| | | |
|-----------------------------------|----------------------------|---|
| | | <p>CL_FP_SOFT_FLOAT – Basic floating-point operations (such as addition, subtraction, multiplication) are implemented in software.</p> <p>The mandated minimum floating-point capability for devices that are not of type CL_DEVICE_TYPE_CUSTOM is: CL_FP_ROUND_TO_NEAREST CL_FP_INF_NAN.</p> |
| <p>CL_DEVICE_DOUBLE_FP_CONFIG</p> | <p>cl_device_fp_config</p> | <p>Describes double precision floating-point capability of the OpenCL device. This is a bit-field that describes one or more of the following values:</p> <p>CL_FP_DENORM – denorms are supported</p> <p>CL_FP_INF_NAN – INF and NaNs are supported.</p> <p>CL_FP_ROUND_TO_NEAREST – round to nearest even rounding mode supported.</p> <p>CL_FP_ROUND_TO_ZERO – round to zero rounding mode supported.</p> <p>CL_FP_ROUND_TO_INF – round to positive and negative infinity rounding modes supported.</p> <p>CP_FP_FMA – IEEE754-2008 fused multiply-add is supported.</p> <p>CL_FP_SOFT_FLOAT – Basic floating-point operations (such as addition, subtraction, multiplication) are implemented in software.</p> <p>Double precision is an optional feature so the mandated minimum double precision floating-point capability is 0.</p> <p>If double precision is supported by the device, then the minimum double precision floating-point capability must be: CL_FP_FMA CL_FP_ROUND_TO_NEAREST CL_FP_ROUND_TO_ZERO CL_FP_ROUND_TO_INF </p> |

| | | |
|---|--------------------------|---|
| | | CL_FP_INF_NAN CL_FP_DENORM. |
| | | |
| CL_DEVICE_GLOBAL_MEM_CACHE_TYPE | cl_device_mem_cache_type | Type of global memory cache supported. Valid values are: CL_NONE, CL_READ_ONLY_CACHE and CL_READ_WRITE_CACHE. |
| CL_DEVICE_GLOBAL_MEM_CACHELINE_SIZE | cl_uint | Size of global memory cache line in bytes. |
| CL_DEVICE_GLOBAL_MEM_CACHE_SIZE | cl_ulong | Size of global memory cache in bytes. |
| CL_DEVICE_GLOBAL_MEM_SIZE | cl_ulong | Size of global device memory in bytes. |
| | | |
| CL_DEVICE_MAX_CONSTANT_BUFFER_SIZE | cl_ulong | Max size in bytes of a constant buffer allocation. The minimum value is 64 KB for devices that are not of type CL_DEVICE_TYPE_CUSTOM. |
| CL_DEVICE_MAX_CONSTANT_ARGS | cl_uint | Max number of arguments declared with the __constant qualifier in a kernel. The minimum value is 8 for devices that are not of type CL_DEVICE_TYPE_CUSTOM. |
| | | |
| CL_DEVICE_LOCAL_MEM_TYPE | cl_device_local_mem_type | Type of local memory supported. This can be set to CL_LOCAL implying dedicated local memory storage such as SRAM, or CL_GLOBAL. For custom devices, CL_NONE can also be returned indicating no local memory support. |
| CL_DEVICE_LOCAL_MEM_SIZE | cl_ulong | Size of local memory arena in bytes. The minimum value is 32 KB for devices that are not of type CL_DEVICE_TYPE_CUSTOM. |
| CL_DEVICE_ERROR_CORRECTION_SUPPORT | cl_bool | Is CL_TRUE if the device implements error correction for all accesses to compute device memory (global and constant). Is CL_FALSE if the device does not implement such error correction. |
| | | |
| CL_DEVICE_HOST_UNIFIED_MEMORY | cl_bool | Is CL_TRUE if the device and the host have a unified memory subsystem and is CL_FALSE otherwise. |
| | | |
| CL_DEVICE_PROFILING_TIMER_RESOLUTION | size_t | Describes the resolution of device timer. This is measured in nanoseconds. Refer to |

| | | |
|---|-----------------------------|---|
| | | <i>section 5.12</i> for details. |
| CL_DEVICE_ENDIAN_LITTLE | cl_bool | Is CL_TRUE if the OpenCL device is a little endian device and CL_FALSE otherwise. |
| CL_DEVICE_AVAILABLE | cl_bool | Is CL_TRUE if the device is available and CL_FALSE if the device is not available. |
| CL_DEVICE_COMPILER_AVAILABLE | cl_bool | Is CL_FALSE if the implementation does not have a compiler available to compile the program source. Is CL_TRUE if the compiler is available. This can be CL_FALSE for the embedded platform profile only. |
| CL_DEVICE_LINKER_AVAILABLE | cl_bool | Is CL_FALSE if the implementation does not have a linker available. Is CL_TRUE if the linker is available. This can be CL_FALSE for the embedded platform profile only. This must be CL_TRUE if CL_DEVICE_COMPILER_AVAILABLE is CL_TRUE. |
| CL_DEVICE_EXECUTION_CAPABILITIES | cl_device_exec_capabilities | Describes the execution capabilities of the device. This is a bit-field that describes one or more of the following values: CL_EXEC_KERNEL – The OpenCL device can execute OpenCL kernels. CL_EXEC_NATIVE_KERNEL – The OpenCL device can execute native kernels. The mandated minimum capability is: CL_EXEC_KERNEL. |
| CL_DEVICE_QUEUE_PROPERTIES | cl_command_queue_properties | Describes the command-queue properties supported by the device. This is a bit-field that describes one or more of the following values: CL_QUEUE_OUT_OF_ORDER_EXEC_MODE_ENABLE |

| | | |
|--------------------------------|----------------|---|
| | | <p>CL_QUEUE_PROFILING_ENABLE</p> <p>These properties are described in <i>table 5.1</i>.</p> <p>The mandated minimum capability is: CL_QUEUE_PROFILING_ENABLE.</p> |
| CL_DEVICE_BUILT_IN_KERNELS | char[] | 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. |
| CL_DEVICE_PLATFORM | cl_platform_id | The platform associated with this device. |
| CL_DEVICE_NAME | char[] | Device name string. |
| CL_DEVICE_VENDOR | char[] | Vendor name string. |
| CL_DRIVER_VERSION | char[] | OpenCL software driver version string in the form <i>major number.minor number</i> |
| CL_DEVICE_PROFILE ⁵ | char[] | <p>OpenCL profile string. Returns the profile name supported by the device. The profile name returned can be one of the following strings:</p> <p>FULL_PROFILE – if the device supports the OpenCL specification (functionality defined as part of the core specification and does not require any extensions to be supported).</p> <p>EMBEDDED_PROFILE - if the device supports the OpenCL embedded profile.</p> |
| CL_DEVICE_VERSION | char[] | <p>OpenCL version string. Returns the OpenCL version supported by the device. This version string has the following format:</p> <p><i>OpenCL<space><major_version.minor_version><space><vendor-specific information></i></p> |

⁵ The platform profile returns the profile that is implemented by the OpenCL framework. If the platform profile returned is FULL_PROFILE, the OpenCL framework will support devices that are FULL_PROFILE and may also support devices that are EMBEDDED_PROFILE. The compiler must be available for all devices i.e. CL_DEVICE_COMPILER_AVAILABLE is CL_TRUE. If the platform profile returned is EMBEDDED_PROFILE, then devices that are only EMBEDDED_PROFILE are supported.

| | | |
|----------------------------|--------|---|
| | | The <i>major_version.minor_version</i> value returned will be 1.2. |
| CL_DEVICE_OPENCL_C_VERSION | char[] | <p>OpenCL C version string. Returns the highest OpenCL C version supported by the compiler for this device that is not of type CL_DEVICE_TYPE_CUSTOM. This version string has the following format:</p> <p><i>OpenCL<space>C<space><major_version.minor_version><space><vendor-specific information></i></p> <p>The <i>major_version.minor_version</i> value returned must be 1.2 if CL_DEVICE_VERSION is OpenCL 1.2.</p> <p>The <i>major_version.minor_version</i> value returned must be 1.1 if CL_DEVICE_VERSION is OpenCL 1.1.</p> <p>The <i>major_version.minor_version</i> value returned can be 1.0 or 1.1 if CL_DEVICE_VERSION is OpenCL 1.0.</p> |
| CL_DEVICE_EXTENSIONS | char[] | <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> <p>cl_khr_int64_base_atomics cl_khr_int64_extended_atomics cl_khr_3d_image_writes cl_khr_fp16 cl_khr_gl_sharing cl_khr_gl_event cl_khr_d3d10_sharing cl_khr_dx9_media_sharing cl_khr_d3d11_sharing cl_khr_depth_images cl_khr_gl_depth_images cl_khr_gl_msaa_sharing cl_khr_image2d_from_buffer cl_khr_initialize_memory cl_khr_context_abort cl_khr_spir</p> |

| | | |
|--|--------------------------------|---|
| | | <p>The following approved Khronos extension names must be returned by all device that support OpenCL C 1.2:</p> <p>cl_khr_global_int32_base_atomics cl_khr_global_int32_extended_atomics cl_khr_local_int32_base_atomics cl_khr_local_int32_extended_atomics cl_khr_byte_addressable_store cl_khr_fp64 (for backward compatibility if double precision is supported)</p> <p>Please refer to the OpenCL 1.2 Extension Specification for a detailed description of these extensions.</p> |
| CL_DEVICE_PRINTF_BUFFER_SIZE | size_t | Maximum size of the internal buffer that holds the output of printf calls from a kernel. The minimum value for the FULL profile is 1 MB. |
| CL_DEVICE_PREFERRED_INTEROP_USER_SYNC | cl_bool | Is CL_TRUE if the device's preference is for the user to be responsible for synchronization, when sharing memory objects between OpenCL and other APIs such as DirectX, CL_FALSE if the device / implementation has a performant path for performing synchronization of memory object shared between OpenCL and other APIs such as DirectX. |
| CL_DEVICE_PARENT_DEVICE | cl_device_id | Returns the cl_device_id of the parent device to which this sub-device belongs. If <i>device</i> is a root-level device, a NULL value is returned. |
| CL_DEVICE_PARTITION_MAX_SUB_DEVICES | cl_uint | Returns the maximum number of sub-devices that can be created when a device is partitioned. The value returned cannot exceed CL_DEVICE_MAX_COMPUTE_UNITS. |
| CL_DEVICE_PARTITION_PROPERTIES | cl_device_partition_property[] | Returns the list of partition types supported by <i>device</i> . The is an array of cl_device_partition_property values drawn from the following list: |

| | | |
|-------------------------------------|--------------------------------|--|
| | | <p>CL_DEVICE_PARTITION_EQUALLY CL_DEVICE_PARTITION_BY_COUNTS CL_DEVICE_PARTITION_BY_AFFINITY_DOMAIN</p> <p>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.</p> |
| CL_DEVICE_PARTITION_AFFINITY_DOMAIN | cl_device_affinity_domain | <p>Returns the list of supported affinity domains for partitioning the <i>device</i> using CL_DEVICE_PARTITION_BY_AFFINITY_DOMAIN. This is a bit-field that describes one or more of the following values:</p> <p>CL_DEVICE_AFFINITY_DOMAIN_NUMA CL_DEVICE_AFFINITY_DOMAIN_L4_CACHE CL_DEVICE_AFFINITY_DOMAIN_L3_CACHE CL_DEVICE_AFFINITY_DOMAIN_L2_CACHE CL_DEVICE_AFFINITY_DOMAIN_L1_CACHE CL_DEVICE_AFFINITY_DOMAIN_NEXT_PARTITIONABLE</p> <p>If the device does not support any affinity domains, a value of 0 will be returned.</p> |
| CL_DEVICE_PARTITION_TYPE | cl_device_partition_property[] | <p>Returns the <i>properties</i> argument specified in clCreateSubDevices if <i>device</i> is a sub-device. In the case where the <i>properties</i> argument to clCreateSubDevices is CL_DEVICE_PARTITION_BY_AFFINITY_DOMAIN, CL_DEVICE_AFFINITY_DOMAIN_NEXT_PARTITIONABLE, the affinity domain used to perform the partition will be returned. This can be one of the following values:</p> <p>CL_DEVICE_AFFINITY_DOMAIN_NUMA CL_DEVICE_AFFINITY_DOMAIN_L4_CACHE CL_DEVICE_AFFINITY_DOMAIN_L3_CACHE CL_DEVICE_AFFINITY_DOMAIN_L2_CACHE CL_DEVICE_AFFINITY_DOMAIN_L1_CACHE</p> <p>Otherwise the implementation may either return a <i>param_value_size_ret</i> 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 <i>param_value</i> points to.</p> |
| | | |
| CL_DEVICE_REFERENCE_COUNT | cl_uint | Returns the <i>device</i> reference count. If the device is a root-level device, a reference |

| | | |
|--|--|---------------------------|
| | | count of one is returned. |
|--|--|---------------------------|

Table 4.3. *OpenCL Device Queries*

The device queries described in *table 4.3* should return the same information for a root-level device i.e. a device returned by **clGetDeviceIDs** and any sub-devices created from this device except for the following queries:

CL_DEVICE_GLOBAL_MEM_CACHE_SIZE
 CL_DEVICE_BUILT_IN_KERNELS
 CL_DEVICE_PARENT_DEVICE
 CL_DEVICE_PARTITION_TYPE
 CL_DEVICE_REFERENCE_COUNT

clGetDeviceInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_DEVICE if *device* is not valid.
- ✚ CL_INVALID_VALUE if *param_name* is not one of the supported values or if size in bytes specified by *param_value_size* is < size of return type as specified in *table 4.3* and *param_value* is not a NULL value or if *param_name* is a value that is available as an extension and the corresponding extension is not supported by the device.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

4.3 Partitioning a Device

The function

```
cl_int clCreateSubDevices (cl_device_id in_device,
                          const cl_device_partition_property *properties,
                          cl_uint num_devices,
                          cl_device_id *out_devices,
                          cl_uint *num_devices_ret)
```

creates an array of sub-devices that each reference a non-intersecting set of compute units within *in_device*, according to a partition scheme given by *properties*. The output sub-devices may be used in every way that the root (or parent) device can be used, including creating contexts, building programs, further calls to **clCreateSubDevices** and creating command-queues. When a command-queue is created against a sub-device, the commands enqueued on the queue are executed only on the sub-device.

in_device is the device to be partitioned.

properties specifies how *in_device* is to be partition described by a partition name and its corresponding value. Each partition name is immediately followed by the corresponding desired value. The list is terminated with 0. The list of supported partitioning schemes is described in *table 4.4*. Only one of the listed partitioning schemes can be specified in *properties*.

| cl_device_partition_property enum | Partition value | Description |
|-----------------------------------|-----------------|---|
| CL_DEVICE_PARTITION_EQUALLY | unsigned int | Split the aggregate device into as many smaller aggregate devices as can be created, each containing <i>n</i> compute units. The value <i>n</i> is passed as the value accompanying this property. If <i>n</i> does not divide evenly into CL_DEVICE_PARTITION_MAX_COMPUTE_UNITS, then the remaining compute units are not used. |
| CL_DEVICE_PARTITION_BY_COUNTS | unsigned int | This property is followed by a CL_DEVICE_PARTITION_BY_COUNTS_LIST_END terminated list of compute unit counts. For each non-zero count <i>m</i> in the list, a sub-device is created with <i>m</i> compute units in it. CL_DEVICE_PARTITION_BY_COUNTS_LIST_END is defined to be 0. The number of non-zero count entries in the list may not exceed CL_DEVICE_PARTITION_MAX_SUB_DEVICES. |

| | | |
|--|---------------------------|---|
| | | <p>The total number of compute units specified may not exceed CL_DEVICE_PARTITION_MAX_COMPUTE_UNITS.</p> |
| CL_DEVICE_PARTITION_BY_AFFINITY_DOMAIN | cl_device_affinity_domain | <p>Split the device into smaller aggregate devices containing one or more compute units that all share part of a cache hierarchy. The value accompanying this property may be drawn from the following list:</p> <p>CL_DEVICE_AFFINITY_DOMAIN_NUMA – Split the device into sub-devices comprised of compute units that share a NUMA node.</p> <p>CL_DEVICE_AFFINITY_DOMAIN_L4_CACHE – Split the device into sub-devices comprised of compute units that share a level 4 data cache.</p> <p>CL_DEVICE_AFFINITY_DOMAIN_L3_CACHE – Split the device into sub-devices comprised of compute units that share a level 3 data cache.</p> <p>CL_DEVICE_AFFINITY_DOMAIN_L2_CACHE – Split the device into sub-devices comprised of compute units that share a level 2 data cache.</p> <p>CL_DEVICE_AFFINITY_DOMAIN_L1_CACHE – Split the device into sub-devices comprised of compute units that share a level 1 data cache.</p> <p>CL_DEVICE_AFFINITY_DOMAIN_NEXT_PARTITIONABLE – Split the device along the next partitionable affinity domain. The implementation shall find the first level along which the device or sub-device may be further subdivided in the order NUMA, L4, L3, L2, L1, and partition the device into sub-devices comprised of compute units that share memory subsystems at this level.</p> <p>The user may determine what happened by calling clGetDeviceInfo(CL_DEVICE_PARTITION_TYPE) on the sub-devices.</p> |

Table 4.4 *List of supported partition schemes by clCreateSubDevices*

num_devices is the size of memory pointed to by *out_devices* specified as the number of *cl_device_id* entries.

out_devices is the buffer where the OpenCL sub-devices will be returned. If *out_devices* is NULL, this argument is ignored. If *out_devices* is not NULL, *num_devices* must be greater than or equal to the number of sub-devices that *device* may be partitioned into according to the partitioning scheme specified in *properties*.

num_devices_ret returns the number of sub-devices that *device* may be partitioned into according to the partitioning scheme specified in *properties*. If *num_devices_ret* is NULL, it is ignored.

clCreateSubDevices returns CL_SUCCESS if the partition is created successfully. Otherwise, it returns a NULL value with the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_DEVICE if *in_device* is not valid.
- ✚ CL_INVALID_VALUE if values specified in *properties* are not valid or if values specified in *properties* are valid but not supported by the device.
- ✚ CL_INVALID_VALUE if *out_devices* is not NULL and *num_devices* is less than the number of sub-devices created by the partition scheme.
- ✚ CL_DEVICE_PARTITION_FAILED if the partition name is supported by the implementation but *in_device* could not be further partitioned.
- ✚ CL_INVALID_DEVICE_PARTITION_COUNT if the partition name specified in *properties* is CL_DEVICE_PARTITION_BY_COUNTS and the number of sub-devices requested exceeds CL_DEVICE_PARTITION_MAX_SUB_DEVICES or the total number of compute units requested exceeds CL_DEVICE_PARTITION_MAX_COMPUTE_UNITS for *in_device*, or the number of compute units requested for one or more sub-devices is less than zero or the number of sub-devices requested exceeds CL_DEVICE_PARTITION_MAX_COMPUTE_UNITS for *in_device*.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

A few examples that describe how to specify partition properties in *properties* argument to **clCreateSubDevices** are given below:

To partition a device containing 16 compute units into two sub-devices, each containing 8 compute units, pass the following in *properties*:

```
{ CL_DEVICE_PARTITION_EQUALLY, 8, 0 }
```

To partition a device with four compute units into two sub-devices with one sub-device containing 3 compute units and the other sub-device 1 compute unit, pass the following in properties argument:

```
{ CL_DEVICE_PARTITION_BY_COUNTS,  
  3, 1, CL_DEVICE_PARTITION_BY_COUNTS_LIST_END, 0 }
```

To split a device along the outermost cache line (if any), pass the following in properties argument:

```
{ CL_DEVICE_PARTITION_BY_AFFINITY_DOMAIN,  
  CL_DEVICE_AFFINITY_DOMAIN_NEXT_PARTITIONABLE,  
  0 }
```

The function

```
cl_int      clRetainDevice (cl_device_id device)
```

increments the *device* reference count if *device* is a valid sub-device created by a call to **clCreateSubDevices**. If *device* is a root level device i.e. a `cl_device_id` returned by **clGetDeviceIDs**, the *device* reference count remains unchanged. **clRetainDevice** returns `CL_SUCCESS` if the function is executed successfully or the device is a root-level device. Otherwise, it returns one of the following errors:

- ✚ `CL_INVALID_DEVICE` if *device* is not a valid sub-device created by a call to **clCreateSubDevices**.
- ✚ `CL_OUT_OF_RESOURCES` if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ `CL_OUT_OF_HOST_MEMORY` if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```
cl_int      clReleaseDevice (cl_device_id device)
```

decrements the *device* reference count if *device* is a valid sub-device created by a call to **clCreateSubDevices**. If *device* is a root level device i.e. a `cl_device_id` returned by **clGetDeviceIDs**, the *device* reference count remains unchanged. **clReleaseDevice** returns `CL_SUCCESS` if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_DEVICE if *device* is not a valid sub-device created by a call to **clCreateSubDevices**.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

After the *device* reference count becomes zero and all the objects attached to *device* (such as command-queues) are released, the *device* object is deleted.

4.4 Contexts

The function

```
cl_context    clCreateContext (const cl_context_properties *properties,
                             cl_uint num_devices,
                             const cl_device_id *devices,
                             void (CL_CALLBACK *pfn_notify)(const char *errinfo,
                                                             const void *private_info, size_t cb,
                                                             void *user_data),
                             void *user_data,
                             cl_int *errcode_ret)
```

creates an OpenCL context. An OpenCL context is created with one or more devices. Contexts are used by the OpenCL runtime for managing objects such as command-queues, memory, program and kernel objects and for executing kernels on one or more devices specified in the context.

properties specifies a list of context property names and their corresponding values. Each property name is immediately followed by the corresponding desired value. The list is terminated with 0. The list of supported properties is described in *table 4.5*. *properties* can be NULL in which case the platform that is selected is implementation-defined.

| cl_context_properties enum | Property value | Description |
|------------------------------|----------------|--|
| CL_CONTEXT_PLATFORM | cl_platform_id | Specifies the platform to use. |
| CL_CONTEXT_INTEROP_USER_SYNC | cl_bool | Specifies whether the user is responsible for synchronization between OpenCL and other APIs. Please refer to the specific sections in the OpenCL 1.2 extension specification that describe sharing with other APIs for restrictions on using this flag. If CL_CONTEXT_INTEROP_USER_SYNC is not specified, a default of CL_FALSE is assumed. |

Table 4.5 List of supported properties by `clCreateContext`

num_devices is the number of devices specified in the *devices* argument.

devices is a pointer to a list of unique devices⁶ returned by **clGetDeviceIDs** or sub-devices created by **clCreateSubDevices** for a platform.

pfm_notify is a callback function that can be registered by the application. This callback function will be used by the OpenCL implementation to report information on errors during context creation as well as errors that occur at runtime in this context. This callback function may be called asynchronously by the OpenCL implementation. It is the application's responsibility to ensure that the callback function is thread-safe. The parameters to this callback function are:

- ✚ *errinfo* is a pointer to an error string.
- ✚ *private_info* and *cb* represent a pointer to binary data that is returned by the OpenCL implementation that can be used to log additional information helpful in debugging the error.
- ✚ *user_data* is a pointer to user supplied data.

If *pfm_notify* is NULL, no callback function is registered.

NOTE: There are a number of cases where error notifications need to be delivered due to an error that occurs outside a context. Such notifications may not be delivered through the *pfm_notify* callback. Where these notifications go is implementation-defined.

user_data will be passed as the *user_data* argument when *pfm_notify* is called. *user_data* can be NULL.

errcode_ret will return an appropriate error code. If *errcode_ret* is NULL, no error code is returned.

clCreateContext returns a valid non-zero context and *errcode_ret* is set to CL_SUCCESS if the context is created successfully. Otherwise, it returns a NULL value with the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_PLATFORM if *properties* is NULL and no platform could be selected or if platform value specified in *properties* is not a valid platform.
- ✚ CL_INVALID_PROPERTY if context property name in *properties* is not a supported property name, if the value specified for a supported property name is not valid, or if the same property name is specified more than once.
- ✚ CL_INVALID_VALUE if *devices* is NULL.
- ✚ CL_INVALID_VALUE if *num_devices* is equal to zero.
- ✚ CL_INVALID_VALUE if *pfm_notify* is NULL but *user_data* is not NULL.
- ✚ CL_INVALID_DEVICE if *devices* contains an invalid device.

⁶ Duplicate devices specified in *devices* are ignored.

- ✚ CL_DEVICE_NOT_AVAILABLE if a device in *devices* is currently not available even though the device was returned by **clGetDeviceIDs**.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```

cl_context
clCreateContextFromType7(const cl_context_properties *properties,
                          cl_device_type device_type,
                          void (CL_CALLBACK *pfn_notify)(const char *errinfo,
                                                         const void *private_info, size_t cb,
                                                         void *user_data),
                          void *user_data,
                          cl_int *errcode_ret)

```

creates an OpenCL context from a device type that identifies the specific device(s) to use. Only devices that are returned by **clGetDeviceIDs** for *device_type* are used to create the context. The context does not reference any sub-devices that may have been created from these devices.

properties specifies a list of context property names and their corresponding values. Each property name is immediately followed by the corresponding desired value. The list of supported properties is described in *table 4.5*. *properties* can also be NULL in which case the platform that is selected is implementation-defined.

device_type is a bit-field that identifies the type of device and is described in *table 4.2* in *section 4.2*.

pfn_notify and *user_data* are described in **clCreateContext**.

errcode_ret will return an appropriate error code. If *errcode_ret* is NULL, no error code is returned.

clCreateContextFromType returns a valid non-zero context and *errcode_ret* is set to CL_SUCCESS if the context is created successfully. Otherwise, it returns a NULL value with the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_PLATFORM if *properties* is NULL and no platform could be selected or if

⁷ **clCreateContextFromType** may return all or a subset of the actual physical devices present in the platform and that match *device_type*.

platform value specified in *properties* is not a valid platform.

- ✚ CL_INVALID_PROPERTY if context property name in *properties* is not a supported property name, if the value specified for a supported property name is not valid, or if the same property name is specified more than once.
- ✚ CL_INVALID_VALUE if *pfm_notify* is NULL but *user_data* is not NULL.
- ✚ CL_INVALID_DEVICE_TYPE if *device_type* is not a valid value.
- ✚ CL_DEVICE_NOT_AVAILABLE if no devices that match *device_type* and property values specified in *properties* are currently available.
- ✚ CL_DEVICE_NOT_FOUND if no devices that match *device_type* and property values specified in *properties* were found.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

cl_int **clRetainContext** (cl_context *context*)

increments the *context* reference count. **clRetainContext** returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_CONTEXT if *context* is not a valid OpenCL context.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

clCreateContext and **clCreateContextFromType** perform an implicit retain. This is very helpful for 3rd party libraries, which typically get a context passed to them by the application. However, it is possible that the application may delete the context without informing the library. Allowing functions to attach to (i.e. retain) and release a context solves the problem of a context being used by a library no longer being valid.

The function

cl_int **clReleaseContext** (cl_context *context*)

decrements the *context* reference count. **clReleaseContext** returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_CONTEXT if *context* is not a valid OpenCL context.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

After the *context* reference count becomes zero and all the objects attached to *context* (such as memory objects, command-queues) are released, the *context* is deleted.

The function

```
cl_int          clGetContextInfo (cl_context context,
                                cl_context_info param_name,
                                size_t param_value_size,
                                void *param_value,
                                size_t *param_value_size_ret)
```

can be used to query information about a context.

context specifies the OpenCL context being queried.

param_name is an enumeration constant that specifies the information to query.

param_value is a pointer to memory where the appropriate result being queried is returned. If *param_value* is NULL, it is ignored.

param_value_size specifies the size in bytes of memory pointed to by *param_value*. This size must be greater than or equal to the size of return type as described in *table 4.6*.

param_value_size_ret returns the actual size in bytes of data being queried by *param_value*. If *param_value_size_ret* is NULL, it is ignored.

The list of supported *param_name* values and the information returned in *param_value* by **clGetContextInfo** is described in *table 4.6*.

| cl_context_info | Return Type | Information returned in param_value |
|--|-------------------------|---|
| CL_CONTEXT_REFERENCE_COUNT ⁸ | cl_uint | Return the <i>context</i> reference count. |
| CL_CONTEXT_NUM_DEVICES | cl_uint | Return the number of devices in <i>context</i> . |
| CL_CONTEXT_DEVICES | cl_device id[] | Return the list of devices in <i>context</i> . |
| CL_CONTEXT_PROPERTIES | cl_context_properties[] | <p>Return the <i>properties</i> argument specified in clCreateContext or clCreateContextFromType.</p> <p>If the <i>properties</i> argument specified in clCreateContext or clCreateContextFromType used to create <i>context</i> is not NULL, the implementation must return the values specified in the <i>properties</i> argument.</p> <p>If the <i>properties</i> argument specified in clCreateContext or clCreateContextFromType used to create <i>context</i> is NULL, the implementation may return either a <i>param_value_size_ret</i> of 0 i.e. there is no context property value to be returned or can return a context property value of 0 (where 0 is used to terminate the context properties list) in the memory that <i>param_value</i> points to.</p> |

Table 4.6 List of supported *param_names* by **clGetContextInfo**

clGetContextInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_CONTEXT if *context* is not a valid context.
- ✚ CL_INVALID_VALUE if *param_name* is not one of the supported values or if size in bytes specified by *param_value_size* is < size of return type as specified in *table 4.6* and *param_value* is not a NULL value.

⁸ The reference count returned should be considered immediately stale. It is unsuitable for general use in applications. This feature is provided for identifying memory leaks.

- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5. The OpenCL Runtime

In this section we describe the API calls that manage OpenCL objects such as command-queues, memory objects, program objects, kernel objects for `__kernel` functions in a program and calls that allow you to enqueue commands to a command-queue such as executing a kernel, reading, or writing a memory object.

5.1 Command Queues

OpenCL objects such as memory, program and kernel objects are created using a context. Operations on these objects are performed using a command-queue. The command-queue can be used to queue a set of operations (referred to as commands) in order. Having multiple command-queues allows applications to queue multiple independent commands without requiring synchronization. Note that this should work as long as these objects are not being shared. Sharing of objects across multiple command-queues will require the application to perform appropriate synchronization. This is described in *Appendix A*.

The function

```
cl_command_queue  clCreateCommandQueue (cl_context context,
                                         cl_device_id device,
                                         cl_command_queue_properties properties,
                                         cl_int *errcode_ret)
```

creates a command-queue on a specific device.

context must be a valid OpenCL context.

| Command-Queue Properties | Description |
|---|---|
| CL_QUEUE_OUT_OF_ORDER_EXEC_MODE_ENABLE | Determines whether the commands queued in the command-queue are executed in-order or out-of-order. If set, the commands in the command-queue are executed out-of-order. Otherwise, commands are executed in-order. For a detailed description about <code>CL_QUEUE_OUT_OF_ORDER_EXEC_MODE_ENABLE</code> , refer to <i>section 5.11</i> . |
| CL_QUEUE_PROFILING_ENABLE | Enable or disable profiling of commands in the command-queue. If set, the profiling of commands is enabled. Otherwise profiling of commands is disabled. |

| | |
|--|--|
| | For a detailed description, refer to <i>section 5.12</i> . |
|--|--|

Table 5.1 *List of supported `cl_command_queue_property` values and description.*

device must be a device associated with *context*. It can either be in the list of devices specified when *context* is created using **clCreateContext** or have the same device type as device type specified when *context* is created using **clCreateContextFromType**.

properties specifies a list of properties for the command-queue. This is a bit-field and is described in *table 5.1*. Only command-queue properties specified in *table 5.1* can be set in *properties*; otherwise the value specified in *properties* is considered to be not valid..

errcode_ret will return an appropriate error code. If *errcode_ret* is NULL, no error code is returned.

clCreateCommandQueue returns a valid non-zero command-queue and *errcode_ret* is set to CL_SUCCESS if the command-queue is created successfully. Otherwise, it returns a NULL value with one of the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_CONTEXT if *context* is not a valid context.
- ✚ CL_INVALID_DEVICE if *device* is not a valid device or is not associated with *context*.
- ✚ CL_INVALID_VALUE if values specified in *properties* are not valid.
- ✚ CL_INVALID_QUEUE_PROPERTIES if values specified in *properties* are valid but are not supported by the device.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

`cl_int` **clRetainCommandQueue** (`cl_command_queue` *command_queue*)

increments the *command_queue* reference count. **clRetainCommandQueue** returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.

- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

clCreateCommandQueue performs an implicit retain. This is very helpful for 3rd party libraries, which typically get a command-queue passed to them by the application. However, it is possible that the application may delete the command-queue without informing the library. Allowing functions to attach to (i.e. retain) and release a command-queue solves the problem of a command-queue being used by a library no longer being valid.

The function

```
cl_int          clReleaseCommandQueue (cl_command_queue command_queue)
```

decrements the *command_queue* reference count. **clReleaseCommandQueue** returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

After the *command_queue* reference count becomes zero and all commands queued to *command_queue* have finished (eg. kernel executions, memory object updates etc.), the command-queue is deleted.

clReleaseCommandQueue performs an implicit flush to issue any previously queued OpenCL commands in *command_queue*.

The function

```
cl_int          clGetCommandQueueInfo (cl_command_queue command_queue,
                                       cl_command_queue_info param_name,
                                       size_t param_value_size,
                                       void *param_value,
                                       size_t *param_value_size_ret)
```

can be used to query information about a command-queue.

command_queue specifies the command-queue being queried.

param_name specifies the information to query.

param_value is a pointer to memory where the appropriate result being queried is returned. If *param_value* is NULL, it is ignored.

param_value_size is used to specify the size in bytes of memory pointed to by *param_value*. This size must be \geq size of return type as described in *table 5.2*. If *param_value* is NULL, it is ignored.

param_value_size_ret returns the actual size in bytes of data being queried by *param_value*. If *param_value_size_ret* is NULL, it is ignored.

The list of supported *param_name* values and the information returned in *param_value* by **clGetCommandQueueInfo** is described in *table 5.2*.

| cl_command_queue_info | Return Type | Information returned in param_value |
|---------------------------------------|-----------------------------|--|
| CL_QUEUE_CONTEXT | cl_context | Return the context specified when the command-queue is created. |
| CL_QUEUE_DEVICE | cl_device_id | Return the device specified when the command-queue is created. |
| CL_QUEUE_REFERENCE_COUNT ⁹ | cl_uint | Return the command-queue reference count. |
| CL_QUEUE_PROPERTIES | cl_command_queue_properties | Return the currently specified properties for the command-queue. These properties are specified by the <i>properties</i> argument in clCreateCommandQueue . |

Table 5.2 List of supported *param_names* by *clGetCommandQueueInfo*

clGetCommandQueueInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_VALUE if *param_name* is not one of the supported values or if size in bytes specified by *param_value_size* is $<$ size of return type as specified in *table 5.2* and *param_value* is not a NULL value.

⁹ The reference count returned should be considered immediately stale. It is unsuitable for general use in applications. This feature is provided for identifying memory leaks.

- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

NOTE

It is possible that a device(s) becomes unavailable after a context and command-queues that use this device(s) have been created and commands have been queued to command-queues. In this case the behavior of OpenCL API calls that use this context (and command-queues) are considered to be implementation-defined. The user callback function, if specified, when the context is created can be used to record appropriate information in the *errinfo*, *private_info* arguments passed to the callback function when the device becomes unavailable.

5.2 Buffer Objects

A *buffer* object stores a one-dimensional collection of elements. Elements of a *buffer* object can be a scalar data type (such as an int, float), vector data type, or a user-defined structure.

5.2.1 Creating Buffer Objects

A **buffer object** is created using the following function

```
cl_mem  clCreateBuffer (cl_context context,
                      cl_mem_flags flags,
                      size_t size,
                      void *host_ptr,
                      cl_int *errcode_ret)
```

context is a valid OpenCL context used to create the buffer object.

flags is a bit-field that is used to specify allocation and usage information such as the memory arena that should be used to allocate the buffer object and how it will be used. *Table 5.3* describes the possible values for *flags*. If value specified for *flags* is 0, the default is used which is CL_MEM_READ_WRITE.

| cl_mem_flags | Description |
|---------------------|--|
| CL_MEM_READ_WRITE | This flag specifies that the memory object will be read and written by a kernel. This is the default. |
| CL_MEM_WRITE_ONLY | This flag specifies that the memory object will be written but not read by a kernel. Reading from a buffer or image object created with CL_MEM_WRITE_ONLY inside a kernel is undefined. CL_MEM_READ_WRITE and CL_MEM_WRITE_ONLY are mutually exclusive. |
| CL_MEM_READ_ONLY | This flag specifies that the memory object is a read-only memory object when used inside a kernel. Writing to a buffer or image object created with CL_MEM_READ_ONLY inside a kernel is undefined. CL_MEM_READ_WRITE or CL_MEM_WRITE_ONLY and CL_MEM_READ_ONLY are mutually exclusive. |
| CL_MEM_USE_HOST_PTR | This flag is valid only if <i>host_ptr</i> is not NULL. If |

| | |
|-------------------------------|---|
| | <p>specified, it indicates that the application wants the OpenCL implementation to use memory referenced by <i>host_ptr</i> as the storage bits for the memory object.</p> <p>OpenCL implementations are allowed to cache the buffer contents pointed to by <i>host_ptr</i> in device memory. This cached copy can be used when kernels are executed on a device.</p> <p>The result of OpenCL commands that operate on multiple buffer objects created with the same <i>host_ptr</i> or overlapping host regions is considered to be undefined.</p> <p>Also refer to <i>section C.3</i> for a description of the alignment rules for <i>host_ptr</i> for memory objects (buffer and images) created using <code>CL_MEM_USE_HOST_PTR</code>.</p> |
| CL_MEM_ALLOC_HOST_PTR | <p>This flag specifies that the application wants the OpenCL implementation to allocate memory from host accessible memory.</p> <p><code>CL_MEM_ALLOC_HOST_PTR</code> and <code>CL_MEM_USE_HOST_PTR</code> are mutually exclusive.</p> |
| CL_MEM_COPY_HOST_PTR | <p>This flag is valid only if <i>host_ptr</i> is not NULL. If specified, it indicates that the application wants the OpenCL implementation to allocate memory for the memory object and copy the data from memory referenced by <i>host_ptr</i>.</p> <p><code>CL_MEM_COPY_HOST_PTR</code> and <code>CL_MEM_USE_HOST_PTR</code> are mutually exclusive.</p> <p><code>CL_MEM_COPY_HOST_PTR</code> can be used with <code>CL_MEM_ALLOC_HOST_PTR</code> to initialize the contents of the <code>cl_mem</code> object allocated using host-accessible (e.g. PCIe) memory.</p> |
| CL_MEM_HOST_WRITE_ONLY | <p>This flag specifies that the host will only write to the memory object (using OpenCL APIs that enqueue a write or a map for write). This can be used to optimize write access from the host (e.g. enable write-combined allocations for memory objects for devices that communicate with the host over a system bus such as PCIe).</p> |

| | |
|------------------------------|---|
| CL_MEM_HOST_READ_ONLY | This flag specifies that the host will only read the memory object (using OpenCL APIs that enqueue a read or a map for read). CL_MEM_HOST_WRITE_ONLY and CL_MEM_HOST_READ_ONLY are mutually exclusive. |
| CL_MEM_HOST_NO_ACCESS | This flag specifies that the host will not read or write the memory object. CL_MEM_HOST_WRITE_ONLY or CL_MEM_HOST_READ_ONLY and CL_MEM_HOST_NO_ACCESS are mutually exclusive. |

Table 5.3 List of supported *cl_mem_flags* values

size is the size in bytes of the buffer memory object to be allocated.

host_ptr is a pointer to the buffer data that may already be allocated by the application. The size of the buffer that *host_ptr* points to must be \geq *size* bytes.

errcode_ret will return an appropriate error code. If *errcode_ret* is NULL, no error code is returned.

clCreateBuffer returns a valid non-zero buffer object and *errcode_ret* is set to CL_SUCCESS if the buffer object is created successfully. Otherwise, it returns a NULL value with one of the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_CONTEXT if *context* is not a valid context.
- ✚ CL_INVALID_VALUE if values specified in *flags* are not valid as defined in *table 5.3*.
- ✚ CL_INVALID_BUFFER_SIZE if *size* is 0^{10} .
- ✚ CL_INVALID_HOST_PTR if *host_ptr* is NULL and CL_MEM_USE_HOST_PTR or CL_MEM_COPY_HOST_PTR are set in *flags* or if *host_ptr* is not NULL but CL_MEM_COPY_HOST_PTR or CL_MEM_USE_HOST_PTR are not set in *flags*.
- ✚ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for buffer object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the

¹⁰ Implementations may return CL_INVALID_BUFFER_SIZE if *size* is greater than CL_DEVICE_MAX_MEM_ALLOC_SIZE value specified in *table 4.3* for all *devices* in context.

OpenCL implementation on the device.

- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```
cl_mem      clCreateSubBuffer (cl_mem buffer,
                               cl_mem_flags flags,
                               cl_buffer_create_type buffer_create_type,
                               const void *buffer_create_info,
                               cl_int *errcode_ret)
```

can be used to create a new buffer object (referred to as a sub-buffer object) from an existing buffer object.

buffer must be a valid buffer object and cannot be a sub-buffer object.

flags is a bit-field that is used to specify allocation and usage information about the sub-buffer memory object being created and is described in *table 5.3*. If the CL_MEM_READ_WRITE, CL_MEM_READ_ONLY or CL_MEM_WRITE_ONLY values are not specified in *flags*, they are inherited from the corresponding memory access qualifiers associated with *buffer*. The CL_MEM_USE_HOST_PTR, CL_MEM_ALLOC_HOST_PTR and CL_MEM_COPY_HOST_PTR values cannot be specified in *flags* but are inherited from the corresponding memory access qualifiers associated with *buffer*. If CL_MEM_COPY_HOST_PTR is specified in the memory access qualifier values associated with *buffer* it does not imply any additional copies when the sub-buffer is created from *buffer*. If the CL_MEM_HOST_WRITE_ONLY, CL_MEM_HOST_READ_ONLY or CL_MEM_HOST_NO_ACCESS values are not specified in *flags*, they are inherited from the corresponding memory access qualifiers associated with *buffer*.

buffer_create_type and *buffer_create_info* describe the type of buffer object to be created. The list of supported values for *buffer_create_type* and corresponding descriptor that *buffer_create_info* points to is described in *table 5.4*.

| cl_buffer_create_type | Description |
|------------------------------|--|
| CL_BUFFER_CREATE_TYPE_REGION | Create a buffer object that represents a specific region in <i>buffer</i> . <i>buffer_create_info</i> is a pointer to the following structure: <pre>typedef struct _cl_buffer_region { size_t origin; size_t size; } cl_buffer_region;</pre> |

| | |
|--|---|
| | <p>(<i>origin</i>, <i>size</i>) defines the offset and size in bytes in <i>buffer</i>.</p> <p>If <i>buffer</i> is created with CL_MEM_USE_HOST_PTR, the <i>host_ptr</i> associated with the buffer object returned is <i>host_ptr</i> + <i>origin</i>.</p> <p>The buffer object returned references the data store allocated for <i>buffer</i> and points to a specific region given by (<i>origin</i>, <i>size</i>) in this data store.</p> <p>CL_INVALID_VALUE is returned in <i>errcode_ret</i> if the region specified by (<i>origin</i>, <i>size</i>) is out of bounds in <i>buffer</i>.</p> <p>CL_INVALID_BUFFER_SIZE if <i>size</i> is 0.</p> <p>CL_MISALIGNED_SUB_BUFFER_OFFSET is returned in <i>errcode_ret</i> if there are no devices in context associated with <i>buffer</i> for which the <i>origin</i> value is aligned to the CL_DEVICE_MEM_BASE_ADDR_ALIGN value.</p> |
|--|---|

Table 5.4 List of supported names and values in *clCreateSubBuffer*.

clCreateSubBuffer returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors in *errcode_ret*:

- ✚ CL_INVALID_MEM_OBJECT if *buffer* is not a valid buffer object or is a sub-buffer object.
- ✚ CL_INVALID_VALUE if *buffer* was created with CL_MEM_WRITE_ONLY and *flags* specifies CL_MEM_READ_WRITE or CL_MEM_READ_ONLY, or if *buffer* was created with CL_MEM_READ_ONLY and *flags* specifies CL_MEM_READ_WRITE or CL_MEM_WRITE_ONLY, or if *flags* specifies CL_MEM_USE_HOST_PTR or CL_MEM_ALLOC_HOST_PTR or CL_MEM_COPY_HOST_PTR.
- ✚ CL_INVALID_VALUE if *buffer* was created with CL_MEM_HOST_WRITE_ONLY and *flags* specify CL_MEM_HOST_READ_ONLY, or if *buffer* was created with CL_MEM_HOST_READ_ONLY and *flags* specify CL_MEM_HOST_WRITE_ONLY, or if *buffer* was created with CL_MEM_HOST_NO_ACCESS and *flags* specify CL_MEM_HOST_READ_ONLY or CL_MEM_HOST_WRITE_ONLY.

- ✚ CL_INVALID_VALUE if value specified in *buffer_create_type* is not valid.
- ✚ CL_INVALID_VALUE if value(s) specified in *buffer_create_info* (for a given *buffer_create_type*) is not valid or if *buffer_create_info* is NULL.
- ✚ CL_INVALID_BUFFER_SIZE if *size* is 0.
- ✚ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for sub-buffer object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

NOTE:

Concurrent reading from, writing to and copying between both a buffer object and its sub-buffer object(s) is undefined. Concurrent reading from, writing to and copying between overlapping sub-buffer objects created with the same buffer object is undefined. Only reading from both a buffer object and its sub-buffer objects or reading from multiple overlapping sub-buffer objects is defined.

5.2.2 Reading, Writing and Copying Buffer Objects

The following functions enqueue commands to read from a buffer object to host memory or write to a buffer object from host memory.

```
cl_int clEnqueueReadBuffer (cl_command_queue command_queue,
                             cl_mem buffer,
                             cl_bool blocking_read,
                             size_t offset,
                             size_t size,
                             void *ptr,
                             cl_uint num_events_in_wait_list,
                             const cl_event *event_wait_list,
                             cl_event *event)
```

```
cl_int clEnqueueWriteBuffer (cl_command_queue command_queue,
                              cl_mem buffer,
                              cl_bool blocking_write,
                              size_t offset,
                              size_t size,
                              const void *ptr,
                              cl_uint num_events_in_wait_list,
                              const cl_event *event_wait_list,
                              cl_event *event)
```

command_queue refers to the command-queue in which the read / write command will be queued. *command_queue* and *buffer* must be created with the same OpenCL context.

buffer refers to a valid buffer object.

blocking_read and *blocking_write* indicate if the read and write operations are *blocking* or *non-blocking*.

If *blocking_read* is CL_TRUE i.e. the read command is blocking, **clEnqueueReadBuffer** does not return until the buffer data has been read and copied into memory pointed to by *ptr*.

If *blocking_read* is CL_FALSE i.e. the read command is non-blocking, **clEnqueueReadBuffer** queues a non-blocking read command and returns. The contents of the buffer that *ptr* points to cannot be used until the read command has completed. The *event* argument returns an event object which can be used to query the execution status of the read command. When the read command has completed, the contents of the buffer that *ptr* points to can be used by the application.

If *blocking_write* is CL_TRUE, the OpenCL implementation copies the data referred to by *ptr* and enqueues the write operation in the command-queue. The memory pointed to by *ptr* can be reused by the application after the **clEnqueueWriteBuffer** call returns.

If *blocking_write* is CL_FALSE, the OpenCL implementation will use *ptr* to perform a non-blocking write. As the write is non-blocking the implementation can return immediately. The memory pointed to by *ptr* cannot be reused by the application after the call returns. The *event* argument returns an event object which can be used to query the execution status of the write command. When the write command has completed, the memory pointed to by *ptr* can then be reused by the application.

offset is the offset in bytes in the buffer object to read from or write to.

size is the size in bytes of data being read or written.

ptr is the pointer to buffer in host memory where data is to be read into or to be written from.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular read / write command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueReadBuffer and **clEnqueueWriteBuffer** return CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_CONTEXT if the context associated with *command_queue* and *buffer* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_MEM_OBJECT if *buffer* is not a valid buffer object.
- ✚ CL_INVALID_VALUE if the region being read or written specified by (*offset*, *size*) is out of bounds or if *ptr* is a NULL value or if *size* is 0.

- ✦ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✦ CL_MISALIGNED_SUB_BUFFER_OFFSET if *buffer* is a sub-buffer object and *offset* specified when the sub-buffer object is created is not aligned to CL_DEVICE_MEM_BASE_ADDR_ALIGN value for device associated with *queue*.
- ✦ CL_EXEC_STATUS_ERROR_FOR_EVENTS_IN_WAIT_LIST if the read and write operations are blocking and the execution status of any of the events in *event_wait_list* is a negative integer value.
- ✦ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with *buffer*.
- ✦ CL_INVALID_OPERATION if **clEnqueueReadBuffer** is called on *buffer* which has been created with CL_MEM_HOST_WRITE_ONLY or CL_MEM_HOST_NO_ACCESS.
- ✦ CL_INVALID_OPERATION if **clEnqueueWriteBuffer** is called on *buffer* which has been created with CL_MEM_HOST_READ_ONLY or CL_MEM_HOST_NO_ACCESS.
- ✦ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✦ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The following functions enqueue commands to read a 2D or 3D rectangular region from a buffer object to host memory or write a 2D or 3D rectangular region to a buffer object from host memory.

```

cl_int  clEnqueueReadBufferRect (cl_command_queue command_queue,
                                cl_mem buffer,
                                cl_bool blocking_read,
                                const size_t *buffer_origin,
                                const size_t *host_origin,
                                const size_t *region,
                                size_t buffer_row_pitch,
                                size_t buffer_slice_pitch,
                                size_t host_row_pitch,
                                size_t host_slice_pitch,
                                void *ptr,
                                cl_uint num_events_in_wait_list,
                                const cl_event *event_wait_list,
                                cl_event *event)

```

```

cl_int  clEnqueueWriteBufferRect (cl_command_queue command_queue,
                                  cl_mem buffer,
                                  cl_bool blocking_write,
                                  const size_t *buffer_origin,
                                  const size_t *host_origin,
                                  const size_t *region,
                                  size_t buffer_row_pitch,
                                  size_t buffer_slice_pitch,
                                  size_t host_row_pitch,
                                  size_t host_slice_pitch,
                                  const void *ptr,
                                  cl_uint num_events_in_wait_list,
                                  const cl_event *event_wait_list,
                                  cl_event *event)

```

command_queue refers to the command-queue in which the read / write command will be queued. *command_queue* and *buffer* must be created with the same OpenCL context.

buffer refers to a valid buffer object.

blocking_read and *blocking_write* indicate if the read and write operations are *blocking* or *non-blocking*.

If *blocking_read* is CL_TRUE i.e. the read command is blocking, **clEnqueueReadBufferRect** does not return until the buffer data has been read and copied into memory pointed to by *ptr*.

If *blocking_read* is CL_FALSE i.e. the read command is non-blocking, **clEnqueueReadBufferRect** queues a non-blocking read command and returns. The contents of the buffer that *ptr* points to cannot be used until the read command has completed. The *event* argument returns an event object which can be used to query the execution status of the read command. When the read command has completed, the contents of the buffer that *ptr* points to can be used by the application.

If *blocking_write* is CL_TRUE, the OpenCL implementation copies the data referred to by *ptr* and enqueues the write operation in the command-queue. The memory pointed to by *ptr* can be reused by the application after the **clEnqueueWriteBufferRect** call returns.

If *blocking_write* is CL_FALSE, the OpenCL implementation will use *ptr* to perform a non-blocking write. As the write is non-blocking the implementation can return immediately. The memory pointed to by *ptr* cannot be reused by the application after the call returns. The *event* argument returns an event object which can be used to query the execution status of the write command. When the write command has completed, the memory pointed to by *ptr* can then be reused by the application.

buffer_origin defines the (x, y, z) offset in the memory region associated with *buffer*. For a 2D rectangle region, the z value given by *buffer_origin*[2] should be 0. The offset in bytes is computed as $buffer_origin[2] * buffer_slice_pitch + buffer_origin[1] * buffer_row_pitch + buffer_origin[0]$.

host_origin defines the (x, y, z) offset in the memory region pointed to by *ptr*. For a 2D rectangle region, the z value given by *host_origin*[2] should be 0. The offset in bytes is computed as $host_origin[2] * host_slice_pitch + host_origin[1] * host_row_pitch + host_origin[0]$.

region defines the (*width* in bytes, *height* in rows, *depth* in slices) of the 2D or 3D rectangle being read or written. For a 2D rectangle copy, the *depth* value given by *region*[2] should be 1. The values in *region* cannot be 0.

buffer_row_pitch is the length of each row in bytes to be used for the memory region associated with *buffer*. If *buffer_row_pitch* is 0, *buffer_row_pitch* is computed as *region*[0].

buffer_slice_pitch is the length of each 2D slice in bytes to be used for the memory region associated with *buffer*. If *buffer_slice_pitch* is 0, *buffer_slice_pitch* is computed as $region[1] * buffer_row_pitch$.

host_row_pitch is the length of each row in bytes to be used for the memory region pointed to by *ptr*. If *host_row_pitch* is 0, *host_row_pitch* is computed as *region*[0].

host_slice_pitch is the length of each 2D slice in bytes to be used for the memory region pointed to by *ptr*. If *host_slice_pitch* is 0, *host_slice_pitch* is computed as $region[1] * host_row_pitch$.

ptr is the pointer to *buffer* in host memory where data is to be read into or to be written from.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular read / write command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueReadBufferRect and **clEnqueueWriteBufferRect** return CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✦ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✦ CL_INVALID_CONTEXT if the context associated with *command_queue* and *buffer* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✦ CL_INVALID_MEM_OBJECT if *buffer* is not a valid buffer object.
- ✦ CL_INVALID_VALUE if the region being read or written specified by (*buffer_origin*, *region*, *buffer_row_pitch*, *buffer_slice_pitch*) is out of bounds.
- ✦ CL_INVALID_VALUE if *ptr* is a NULL value.
- ✦ CL_INVALID_VALUE if any *region* array element is 0.
- ✦ CL_INVALID_VALUE if *buffer_row_pitch* is not 0 and is less than *region*[0].
- ✦ CL_INVALID_VALUE if *host_row_pitch* is not 0 and is less than *region*[0].
- ✦ CL_INVALID_VALUE if *buffer_slice_pitch* is not 0 and is less than *region*[1] * *buffer_row_pitch* and not a multiple of *buffer_row_pitch*.
- ✦ CL_INVALID_VALUE if *host_slice_pitch* is not 0 and is less than *region*[1] * *host_row_pitch* and not a multiple of *host_row_pitch*.
- ✦ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✦ CL_MISALIGNED_SUB_BUFFER_OFFSET if *buffer* is a sub-buffer object and *offset* specified when the sub-buffer object is created is not aligned to CL_DEVICE_MEM_BASE_ADDR_ALIGN value for device associated with *queue*.
- ✦ CL_EXEC_STATUS_ERROR_FOR_EVENTS_IN_WAIT_LIST if the read and write operations are blocking and the execution status of any of the events in *event_wait_list* is a negative integer value.
- ✦ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with *buffer*.
- ✦ CL_INVALID_OPERATION if **clEnqueueReadBufferRect** is called on *buffer* which has been created with CL_MEM_HOST_WRITE_ONLY or CL_MEM_HOST_NO_ACCESS.
- ✦ CL_INVALID_OPERATION if **clEnqueueWriteBufferRect** is called on *buffer* which has been created with CL_MEM_HOST_READ_ONLY or CL_MEM_HOST_NO_ACCESS.
- ✦ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the

OpenCL implementation on the device.

- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

NOTE:

Calling **clEnqueueReadBuffer** to read a region of the buffer object with the *ptr* argument value set to *host_ptr + offset*, where *host_ptr* is a pointer to the memory region specified when the buffer object being read is created with CL_MEM_USE_HOST_PTR, must meet the following requirements in order to avoid undefined behavior:

- All commands that use this buffer object or a memory object (buffer or image) created from this buffer object have finished execution before the read command begins execution.
- The buffer object or memory objects created from this buffer object are not mapped.
- The buffer object or memory objects created from this buffer object are not used by any command-queue until the read command has finished execution.

Calling **clEnqueueReadBufferRect** to read a region of the buffer object with the *ptr* argument value set to *host_ptr* and *host_origin*, *buffer_origin* values are the same, where *host_ptr* is a pointer to the memory region specified when the buffer object being read is created with CL_MEM_USE_HOST_PTR, must meet the same requirements given above for **clEnqueueReadBuffer**.

Calling **clEnqueueWriteBuffer** to update the latest bits in a region of the buffer object with the *ptr* argument value set to *host_ptr + offset*, where *host_ptr* is a pointer to the memory region specified when the buffer object being written is created with CL_MEM_USE_HOST_PTR, must meet the following requirements in order to avoid undefined behavior:

- The host memory region given by (*host_ptr + offset*, *cb*) contains the latest bits when the enqueued write command begins execution.
- The buffer object or memory objects created from this buffer object are not mapped.
- The buffer object or memory objects created from this buffer object are not used by any command-queue until the write command has finished execution.

Calling **clEnqueueWriteBufferRect** to update the latest bits in a region of the buffer object with the *ptr* argument value set to *host_ptr* and *host_origin*, *buffer_origin* values are the same, where *host_ptr* is a pointer to the memory region specified when the buffer object being written is created with CL_MEM_USE_HOST_PTR, must meet the following requirements in order to avoid undefined behavior:

- The host memory region given by (*buffer_origin region*) contains the latest bits when the enqueued write command begins execution.
- The buffer object or memory objects created from this buffer object are not mapped.
- The buffer object or memory objects created from this buffer object are not used by any command-queue until the write command has finished execution.

The function

```

cl_int  clEnqueueCopyBuffer (cl_command_queue command_queue,
                             cl_mem src_buffer,
                             cl_mem dst_buffer,
                             size_t src_offset,
                             size_t dst_offset,
                             size_t size,
                             cl_uint num_events_in_wait_list,
                             const cl_event *event_wait_list,
                             cl_event *event)

```

enqueues a command to copy a buffer object identified by *src_buffer* to another buffer object identified by *dst_buffer*.

command_queue refers to the command-queue in which the copy command will be queued. The OpenCL context associated with *command_queue*, *src_buffer* and *dst_buffer* must be the same.

src_offset refers to the offset where to begin copying data from *src_buffer*.

dst_offset refers to the offset where to begin copying data into *dst_buffer*.

size refers to the size in bytes to copy.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular copy command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a

wait for this command to complete. **clEnqueueBarrierWithWaitList** can be used instead. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueCopyBuffer returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_CONTEXT if the context associated with *command_queue*, *src_buffer* and *dst_buffer* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_MEM_OBJECT if *src_buffer* and *dst_buffer* are not valid buffer objects.
- ✚ CL_INVALID_VALUE if *src_offset*, *dst_offset*, *size*, *src_offset + size* or *dst_offset + size* require accessing elements outside the *src_buffer* and *dst_buffer* buffer objects respectively.
- ✚ CL_INVALID_VALUE if *size* is 0.
- ✚ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✚ CL_MISALIGNED_SUB_BUFFER_OFFSET if *src_buffer* is a sub-buffer object and *offset* specified when the sub-buffer object is created is not aligned to CL_DEVICE_MEM_BASE_ADDR_ALIGN value for device associated with *queue*.
- ✚ CL_MISALIGNED_SUB_BUFFER_OFFSET if *dst_buffer* is a sub-buffer object and *offset* specified when the sub-buffer object is created is not aligned to CL_DEVICE_MEM_BASE_ADDR_ALIGN value for device associated with *queue*.
- ✚ CL_MEM_COPY_OVERLAP if *src_buffer* and *dst_buffer* are the same buffer or sub-buffer object and the source and destination regions overlap or if *src_buffer* and *dst_buffer* are different sub-buffers of the same associated buffer object and they overlap. The regions overlap if $src_offset \leq dst_offset \leq src_offset + size - 1$ or if $dst_offset \leq src_offset \leq dst_offset + size - 1$.
- ✚ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with *src_buffer* or *dst_buffer*.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```
cl_int  clEnqueueCopyBufferRect (cl_command_queue command_queue,
                                cl_mem src_buffer,
                                cl_mem dst_buffer,
                                const size_t *src_origin,
                                const size_t *dst_origin,
                                const size_t *region,
                                size_t src_row_pitch,
                                size_t src_slice_pitch,
                                size_t dst_row_pitch,
                                size_t dst_slice_pitch,
                                cl_uint num_events_in_wait_list,
                                const cl_event *event_wait_list,
                                cl_event *event)
```

enqueues a command to copy a 2D or 3D rectangular region from the buffer object identified by *src_buffer* to a 2D or 3D region in the buffer object identified by *dst_buffer*. Copying begins at the source offset and destination offset which are computed as described below in the description for *src_origin* and *dst_origin*. Each byte of the region's width is copied from the source offset to the destination offset. After copying each width, the source and destination offsets are incremented by their respective source and destination row pitches. After copying each 2D rectangle, the source and destination offsets are incremented by their respective source and destination slice pitches.

NOTE: If *src_buffer* and *dst_buffer* are the same buffer object, *src_row_pitch* must equal *dst_row_pitch* and *src_slice_pitch* must equal *dst_slice_pitch*.

command_queue refers to the command-queue in which the copy command will be queued. The OpenCL context associated with *command_queue*, *src_buffer* and *dst_buffer* must be the same.

src_origin defines the (x, y, z) offset in the memory region associated with *src_buffer*. For a 2D rectangle region, the z value given by *src_origin*[2] should be 0. The offset in bytes is computed as *src_origin*[2] * *src_slice_pitch* + *src_origin*[1] * *src_row_pitch* + *src_origin*[0].

dst_origin defines the (x, y, z) offset in the memory region associated with *dst_buffer*. For a 2D rectangle region, the z value given by *dst_origin*[2] should be 0. The offset in bytes is computed as *dst_origin*[2] * *dst_slice_pitch* + *dst_origin*[1] * *dst_row_pitch* + *dst_origin*[0].

region defines the (width in bytes, height in rows, depth in slices) of the 2D or 3D rectangle being copied. For a 2D rectangle, the *depth* value given by *region*[2] should be 1. The values in *region* cannot be 0.

src_row_pitch is the length of each row in bytes to be used for the memory region associated with *src_buffer*. If *src_row_pitch* is 0, *src_row_pitch* is computed as *region*[0].

src_slice_pitch is the length of each 2D slice in bytes to be used for the memory region associated with *src_buffer*. If *src_slice_pitch* is 0, *src_slice_pitch* is computed as *region[1] * src_row_pitch*.

dst_row_pitch is the length of each row in bytes to be used for the memory region associated with *dst_buffer*. If *dst_row_pitch* is 0, *dst_row_pitch* is computed as *region[0]*.

dst_slice_pitch is the length of each 2D slice in bytes to be used for the memory region associated with *dst_buffer*. If *dst_slice_pitch* is 0, *dst_slice_pitch* is computed as *region[1] * dst_row_pitch*.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular copy command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. **clEnqueueBarrierWithWaitList** can be used instead. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueCopyBufferRect returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_CONTEXT if the context associated with *command_queue*, *src_buffer* and *dst_buffer* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_MEM_OBJECT if *src_buffer* and *dst_buffer* are not valid buffer objects.
- ✚ CL_INVALID_VALUE if (*src_origin*, *region*, *src_row_pitch*, *src_slice_pitch*) or (*dst_origin*, *region*, *dst_row_pitch*, *dst_slice_pitch*) require accessing elements outside the *src_buffer* and *dst_buffer* buffer objects respectively.
- ✚ CL_INVALID_VALUE if any *region* array element is 0.
- ✚ CL_INVALID_VALUE if *src_row_pitch* is not 0 and is less than *region[0]*.

- ✦ CL_INVALID_VALUE if *dst_row_pitch* is not 0 and is less than *region[0]*.
- ✦ CL_INVALID_VALUE if *src_slice_pitch* is not 0 and is less than *region[1] * src_row_pitch* or if *src_slice_pitch* is not 0 and is not a multiple of *src_row_pitch*.
- ✦ CL_INVALID_VALUE if *dst_slice_pitch* is not 0 and is less than *region[1] * dst_row_pitch* or if *dst_slice_pitch* is not 0 and is not a multiple of *dst_row_pitch*.
- ✦ CL_INVALID_VALUE if *src_buffer* and *dst_buffer* are the same buffer object and *src_slice_pitch* is not equal to *dst_slice_pitch* and *src_row_pitch* is not equal to *dst_row_pitch*.
- ✦ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✦ CL_MEM_COPY_OVERLAP if *src_buffer* and *dst_buffer* are the same buffer or sub-buffer object and the source and destination regions overlap or if *src_buffer* and *dst_buffer* are different sub-buffers of the same associated buffer object and they overlap. Refer to Appendix E for details on how to determine if source and destination regions overlap.
- ✦ CL_MISALIGNED_SUB_BUFFER_OFFSET if *src_buffer* is a sub-buffer object and *offset* specified when the sub-buffer object is created is not aligned to CL_DEVICE_MEM_BASE_ADDR_ALIGN value for device associated with *queue*.
- ✦ CL_MISALIGNED_SUB_BUFFER_OFFSET if *dst_buffer* is a sub-buffer object and *offset* specified when the sub-buffer object is created is not aligned to CL_DEVICE_MEM_BASE_ADDR_ALIGN value for device associated with *queue*.
- ✦ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with *src_buffer* or *dst_buffer*.
- ✦ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✦ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.2.3 Filling Buffer Objects

The function

```
cl_int  clEnqueueFillBuffer (cl_command_queue command_queue,
                             cl_mem buffer,
                             const void *pattern,
                             size_t pattern_size,
                             size_t offset,
                             size_t size,
                             cl_uint num_events_in_wait_list,
                             const cl_event *event_wait_list,
                             cl_event *event)
```

enqueues a command to fill a buffer object with a pattern of a given pattern size. The usage information which indicates whether the memory object can be read or written by a kernel and/or the host and is given by the `cl_mem_flags` argument value specified when *buffer* is created is ignored by **clEnqueueFillBuffer**.

command_queue refers to the command-queue in which the fill command will be queued. The OpenCL context associated with *command_queue* and *buffer* must be the same.

buffer is a valid buffer object.

pattern is a pointer to the data pattern of size *pattern_size* in bytes. *pattern* will be used to fill a region in *buffer* starting at *offset* and is *size* bytes in size. The data pattern must be a scalar or vector integer or floating-point data type supported by OpenCL as described in sections 6.1.1 and 6.1.2. For example, if *buffer* is to be filled with a pattern of `float4` values, then *pattern* will be a pointer to a `cl_float4` value and *pattern_size* will be `sizeof(cl_float4)`. The maximum value of *pattern_size* is the size of the largest integer or floating-point vector data type supported by the OpenCL device. The memory associated with *pattern* can be reused or freed after the function returns.

offset is the location in bytes of the region being filled in *buffer* and must be a multiple of *pattern_size*.

size is the size in bytes of region being filled in *buffer* and must be a multiple of *pattern_size*.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. **clEnqueueBarrierWithWaitList** can be used instead. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueFillBuffer returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_CONTEXT if the context associated with *command_queue* and *buffer* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_MEM_OBJECT if *buffer* is not a valid buffer object.
- ✚ CL_INVALID_VALUE if *offset* or *offset + size* require accessing elements outside the *buffer* buffer object respectively.
- ✚ CL_INVALID_VALUE if *pattern* is NULL or if *pattern_size* is 0 or if *pattern_size* is not one of {1, 2, 4, 8, 16, 32, 64, 128}.
- ✚ CL_INVALID_VALUE if *offset* and *size* are not a multiple of *pattern_size*.
- ✚ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✚ CL_MISALIGNED_SUB_BUFFER_OFFSET if *buffer* is a sub-buffer object and offset specified when the sub-buffer object is created is not aligned to CL_DEVICE_MEM_BASE_ADDR_ALIGN value for device associated with *queue*.
- ✚ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with *buffer*.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.2.4 Mapping Buffer Objects

The function

```
void * clEnqueueMapBuffer (cl_command_queue command_queue,
                           cl_mem buffer,
                           cl_bool blocking_map,
                           cl_map_flags map_flags,
                           size_t offset,
                           size_t size,
                           cl_uint num_events_in_wait_list,
                           const cl_event *event_wait_list,
                           cl_event *event,
                           cl_int *errcode_ret)
```

enqueues a command to map a region of the buffer object given by *buffer* into the host address space and returns a pointer to this mapped region.

command_queue must be a valid command-queue.

blocking_map indicates if the map operation is *blocking* or *non-blocking*.

If *blocking_map* is CL_TRUE, **clEnqueueMapBuffer** does not return until the specified region in *buffer* is mapped into the host address space and the application can access the contents of the mapped region using the pointer returned by **clEnqueueMapBuffer**.

If *blocking_map* is CL_FALSE i.e. map operation is non-blocking, the pointer to the mapped region returned by **clEnqueueMapBuffer** cannot be used until the map command has completed. The *event* argument returns an event object which can be used to query the execution status of the map command. When the map command is completed, the application can access the contents of the mapped region using the pointer returned by **clEnqueueMapBuffer**.

map_flags is a bit-field and is described in *table 5.5*.

buffer is a valid buffer object. The OpenCL context associated with *command_queue* and *buffer* must be the same.

offset and *size* are the offset in bytes and the size of the region in the buffer object that is being mapped.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must

be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

errcode_ret will return an appropriate error code. If *errcode_ret* is NULL, no error code is returned.

clEnqueueMapBuffer will return a pointer to the mapped region. The *errcode_ret* is set to CL_SUCCESS.

A NULL pointer is returned otherwise with one of the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_CONTEXT if context associated with *command_queue* and *buffer* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_MEM_OBJECT if *buffer* is not a valid buffer object.
- ✚ CL_INVALID_VALUE if region being mapped given by (*offset*, *size*) is out of bounds or if *size* is 0 or if values specified in *map_flags* are not valid.
- ✚ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✚ CL_MISALIGNED_SUB_BUFFER_OFFSET if *buffer* is a sub-buffer object and *offset* specified when the sub-buffer object is created is not aligned to CL_DEVICE_MEM_BASE_ADDR_ALIGN value for device associated with *queue*.
- ✚ CL_MAP_FAILURE if there is a failure to map the requested region into the host address space. This error cannot occur for buffer objects created with CL_MEM_USE_HOST_PTR or CL_MEM_ALLOC_HOST_PTR.
- ✚ CL_EXEC_STATUS_ERROR_FOR_EVENTS_IN_WAIT_LIST if the map operation is blocking and the execution status of any of the events in *event_wait_list* is a negative integer value.

- ✦ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with *buffer*.
- ✦ CL_INVALID_OPERATION if *buffer* has been created with CL_MEM_HOST_WRITE_ONLY or CL_MEM_HOST_NO_ACCESS and CL_MAP_READ is set in *map_flags* or if *buffer* has been created with CL_MEM_HOST_READ_ONLY or CL_MEM_HOST_NO_ACCESS and CL_MAP_WRITE or CL_MAP_WRITE_INVALIDATE_REGION is set in *map_flags*.
- ✦ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✦ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The pointer returned maps a region starting at *offset* and is at least *size* bytes in size. The result of a memory access outside this region is undefined.

NOTE:

If the buffer object is created with CL_MEM_USE_HOST_PTR set in *mem_flags*, the following will be true:

- ✦ The *host_ptr* specified in **clCreateBuffer** is guaranteed to contain the latest bits in the region being mapped when the **clEnqueueMapBuffer** command has completed.
- ✦ The pointer value returned by **clEnqueueMapBuffer** will be derived from the *host_ptr* specified when the buffer object is created.

Mapped buffer objects are unmapped using **clEnqueueUnmapMemObject**. This is described in *section 5.4.2*.

| cl_map_flags | Description |
|--------------|---|
| CL_MAP_READ | <p>This flag specifies that the region being mapped in the memory object is being mapped for reading.</p> <p>The pointer returned by clEnqueueMap{Buffer Image} is guaranteed to contain the latest bits in the region being mapped when the clEnqueueMap{Buffer Image} command has completed</p> |
| CL_MAP_WRITE | <p>This flag specifies that the region being mapped in the memory object is being mapped for writing.</p> |

| | |
|---------------------------------------|---|
| | <p>The pointer returned by clEnqueueMap{Buffer Image} is guaranteed to contain the latest bits in the region being mapped when the clEnqueueMap{Buffer Image} command has completed</p> |
| CL_MAP_WRITE_INVALIDATE_REGION | <p>This flag specifies that the region being mapped in the memory object is being mapped for writing.</p> <p>The contents of the region being mapped are to be discarded. This is typically the case when the region being mapped is overwritten by the host. This flag allows the implementation to no longer guarantee that the pointer returned by clEnqueueMap{Buffer Image} contains the latest bits in the region being mapped which can be a significant performance enhancement.</p> <p>CL_MAP_READ or CL_MAP_WRITE and CL_MAP_WRITE_INVALIDATE_REGION are mutually exclusive.</p> |

Table 5.5 *List of supported cl_map_flags values*

5.3 Image Objects

An *image* object is used to store a one-, two- or three- dimensional texture, frame-buffer or image. The elements of an image object are selected from a list of predefined image formats. The minimum number of elements in a memory object is one.

5.3.1 Creating Image Objects

A **1D image, 1D image buffer, 1D image array, 2D image, 2D image array and 3D image object** can be created using the following function

```
cl_mem  clCreateImage (cl_context context,
                      cl_mem_flags flags,
                      const cl_image_format *image_format,
                      const cl_image_desc *image_desc,
                      void *host_ptr,
                      cl_int *errcode_ret)
```

context is a valid OpenCL context on which the image object is to be created.

flags is a bit-field that is used to specify allocation and usage information about the image memory object being created and is described in *table 5.3*.

For all image types except CL_MEM_OBJECT_IMAGE1D_BUFFER, if value specified for *flags* is 0, the default is used which is CL_MEM_READ_WRITE.

For CL_MEM_OBJECT_IMAGE1D_BUFFER image type, if the CL_MEM_READ_WRITE, CL_MEM_READ_ONLY or CL_MEM_WRITE_ONLY values are not specified in *flags*, they are inherited from the corresponding memory access qualifiers associated with *buffer*. The CL_MEM_USE_HOST_PTR, CL_MEM_ALLOC_HOST_PTR and CL_MEM_COPY_HOST_PTR values cannot be specified in *flags* but are inherited from the corresponding memory access qualifiers associated with *buffer*. If CL_MEM_COPY_HOST_PTR is specified in the memory access qualifier values associated with *buffer* it does not imply any additional copies when the sub-buffer is created from *buffer*. If the CL_MEM_HOST_WRITE_ONLY, CL_MEM_HOST_READ_ONLY or CL_MEM_HOST_NO_ACCESS values are not specified in *flags*, they are inherited from the corresponding memory access qualifiers associated with *buffer*.

image_format is a pointer to a structure that describes format properties of the image to be allocated. Refer to *section 5.3.1.1* for a detailed description of the image format descriptor.

image_desc is a pointer to a structure that describes type and dimensions of the image to be allocated. Refer to *section 5.3.1.2* for a detailed description of the image descriptor.

host_ptr is a pointer to the image data that may already be allocated by the application. Refer to table below for a description of how large the buffer that *host_ptr* points to must be.

| Image Type | Size of buffer that <i>host_ptr</i> points to |
|------------------------------|---|
| CL_MEM_OBJECT_IMAGE1D | \geq image row pitch |
| CL_MEM_OBJECT_IMAGE1D_BUFFER | \geq image row pitch |
| CL_MEM_OBJECT_IMAGE2D | \geq image row pitch * image height |
| CL_MEM_OBJECT_IMAGE3D | \geq image slice pitch * image depth |
| CL_MEM_OBJECT_IMAGE1D_ARRAY | \geq image slice pitch * image array size |
| CL_MEM_OBJECT_IMAGE2D_ARRAY | \geq image slice pitch * image array size |

For a 3D image or 2D image array, the image data specified by *host_ptr* is stored as a linear sequence of adjacent 2D image slices or 2D images respectively. Each 2D image is a linear sequence of adjacent scanlines. Each scanline is a linear sequence of image elements.

For a 2D image, the image data specified by *host_ptr* is stored as a linear sequence of adjacent scanlines. Each scanline is a linear sequence of image elements.

For a 1D image array, the image data specified by *host_ptr* is stored as a linear sequence of adjacent 1D images respectively. Each 1D image or 1D image buffer is a single scanline which is a linear sequence of adjacent elements.

errcode_ret will return an appropriate error code. If *errcode_ret* is NULL, no error code is returned.

clCreateImage returns a valid non-zero image object created and the *errcode_ret* is set to CL_SUCCESS if the image object is created successfully. Otherwise, it returns a NULL value with one of the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_CONTEXT if *context* is not a valid context.
- ✚ CL_INVALID_VALUE if values specified in *flags* are not valid.
- ✚ CL_INVALID_IMAGE_FORMAT_DESCRIPTOR if values specified in *image_format* are not valid or if *image_format* is NULL.
- ✚ CL_INVALID_IMAGE_DESCRIPTOR if values specified in *image_desc* are not valid or if *image_desc* is NULL.
- ✚ CL_INVALID_IMAGE_SIZE if image dimensions specified in *image_desc* exceed the minimum maximum image dimensions described in *table 4.3* for all devices in *context*.
- ✚ CL_INVALID_HOST_PTR if *host_ptr* is NULL and CL_MEM_USE_HOST_PTR or CL_MEM_COPY_HOST_PTR are set in *flags* or if *host_ptr* is not NULL but CL_MEM_COPY_HOST_PTR or CL_MEM_USE_HOST_PTR are not set in *flags*.

- ✚ CL_INVALID_VALUE if a 1D image buffer is being created and the buffer object was created with CL_MEM_WRITE_ONLY and *flags* specifies CL_MEM_READ_WRITE or CL_MEM_READ_ONLY, or if the buffer object was created with CL_MEM_READ_ONLY and *flags* specifies CL_MEM_READ_WRITE or CL_MEM_WRITE_ONLY, or if *flags* specifies CL_MEM_USE_HOST_PTR or CL_MEM_ALLOC_HOST_PTR or CL_MEM_COPY_HOST_PTR.
- ✚ CL_INVALID_VALUE if a 1D image buffer is being created and the buffer object was created with CL_MEM_HOST_WRITE_ONLY and *flags* specifies CL_MEM_HOST_READ_ONLY, or if the buffer object was created with CL_MEM_HOST_READ_ONLY and *flags* specifies CL_MEM_HOST_WRITE_ONLY, or if the buffer object was created with CL_MEM_HOST_NO_ACCESS and *flags* specifies CL_MEM_HOST_READ_ONLY or CL_MEM_HOST_WRITE_ONLY.
- ✚ CL_IMAGE_FORMAT_NOT_SUPPORTED if the *image_format* is not supported.
- ✚ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for image object.
- ✚ CL_INVALID_OPERATION if there are no devices in *context* that support images (i.e. CL_DEVICE_IMAGE_SUPPORT specified in *table 4.3* is CL_FALSE).
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.3.1.1 Image Format Descriptor

The image format descriptor structure is defined as

```
typedef struct _cl_image_format {
    cl_channel_order    image_channel_order;
    cl_channel_type     image_channel_data_type;
} cl_image_format;
```

image_channel_order specifies the number of channels and the channel layout i.e. the memory layout in which channels are stored in the image. Valid values are described in *table 5.6*.

image_channel_data_type describes the size of the channel data type. The list of supported values is described in *table 5.7*. The number of bits per element determined by the *image_channel_data_type* and *image_channel_order* must be a power of two.

| Enum values that can be specified in channel_order |
|---|
| CL_R, CL_Rx or CL_A |
| CL_INTENSITY. This format can only be used if channel data type = CL_UNORM_INT8, CL_UNORM_INT16, CL_SNORM_INT8, CL_SNORM_INT16, CL_HALF_FLOAT or CL_FLOAT. |
| CL_LUMINANCE. This format can only be used if channel data type = CL_UNORM_INT8, CL_UNORM_INT16, CL_SNORM_INT8, CL_SNORM_INT16, CL_HALF_FLOAT or CL_FLOAT. |
| CL_RG, CL_RGx or CL_RA |
| CL_RGB or CL_RGBx. This format can only be used if channel data type = CL_UNORM_SHORT_565, CL_UNORM_SHORT_555, or CL_UNORM_INT_101010. |
| CL_RGBA |
| CL_ARGB, CL_BGRA. This format can only be used if channel data type = CL_UNORM_INT8, CL_SNORM_INT8, CL_SIGNED_INT8 or CL_UNSIGNED_INT8. |


Table 5.6 *List of supported Image Channel Order Values*

| Image Channel Data Type | Description |
|--------------------------------|---|
| CL_SNORM_INT8 | Each channel component is a normalized signed 8-bit integer value |
| CL_SNORM_INT16 | Each channel component is a normalized signed 16-bit integer value |
| CL_UNORM_INT8 | Each channel component is a normalized unsigned 8-bit integer value |
| CL_UNORM_INT16 | Each channel component is a normalized unsigned 16-bit integer value |
| CL_UNORM_SHORT_565 | Represents a normalized 5-6-5 3-channel RGB image. The channel order must be CL_RGB or CL_RGBx. |
| CL_UNORM_SHORT_555 | Represents a normalized x-5-5-5 4-channel xRGB image. The channel order must be CL_RGB or CL_RGBx. |
| CL_UNORM_INT_101010 | Represents a normalized x-10-10-10 4-channel xRGB image. The channel order must be CL_RGB or CL_RGBx. |
| CL_SIGNED_INT8 | Each channel component is an unnormalized signed 8-bit integer value |
| CL_SIGNED_INT16 | Each channel component is an unnormalized signed 16-bit integer value |
| CL_SIGNED_INT32 | Each channel component is an unnormalized signed 32-bit integer value |
| CL_UNSIGNED_INT8 | Each channel component is an unnormalized unsigned 8-bit integer value |


| | |
|--------------------------|---|
| CL_UNSIGNED_INT16 | Each channel component is an unnormalized unsigned 16-bit integer value |
| CL_UNSIGNED_INT32 | Each channel component is an unnormalized unsigned 32-bit integer value |
| CL_HALF_FLOAT | Each channel component is a 16-bit half-float value |
| CL_FLOAT | Each channel component is a single precision floating-point value |

Table 5.7 List of supported Image Channel Data Types

For example, to specify a normalized unsigned 8-bit / channel RGBA image, `image_channel_order = CL_RGBA`, and `image_channel_data_type = CL_UNORM_INT8`. The memory layout of this image format is described below:



Similar, if `image_channel_order = CL_RGBA` and `image_channel_data_type = CL_SIGNED_INT16`, the memory layout of this image format is described below:



`image_channel_data_type` values of `CL_UNORM_SHORT_565`, `CL_UNORM_SHORT_555` and `CL_UNORM_INT_101010` are special cases of packed image formats where the channels of each element are packed into a single unsigned short or unsigned int. For these special packed image formats, the channels are normally packed with the first channel in the most significant bits of the bitfield, and successive channels occupying progressively less significant locations. For `CL_UNORM_SHORT_565`, R is in bits 15:11, G is in bits 10:5 and B is in bits 4:0. For `CL_UNORM_SHORT_555`, bit 15 is undefined, R is in bits 14:10, G in bits 9:5 and B in bits 4:0. For `CL_UNORM_INT_101010`, bits 31:30 are undefined, R is in bits 29:20, G in bits 19:10 and B in bits 9:0.

OpenCL implementations must maintain the minimum precision specified by the number of bits in `image_channel_data_type`. If the image format specified by `image_channel_order`, and `image_channel_data_type` cannot be supported by the OpenCL implementation, then the call to `clCreateImage` will return a NULL memory object.

5.3.1.2 Image Descriptor

The image descriptor structure describes the type and dimensions of the image or image array and is defined as:

```
typedef struct _cl_image_desc {
    cl_mem_object_type  image_type,
    size_t              image_width;
    size_t              image_height;
    size_t              image_depth;
    size_t              image_array_size;
    size_t              image_row_pitch;
    size_t              image_slice_pitch;
    cl_uint             num_mip_levels;
    cl_uint             num_samples;
    cl_mem              buffer;
} cl_image_desc;
```

`image_type` describes the image type and must be either `CL_MEM_OBJECT_IMAGE1D`, `CL_MEM_OBJECT_IMAGE1D_BUFFER`, `CL_MEM_OBJECT_IMAGE1D_ARRAY`, `CL_MEM_OBJECT_IMAGE2D`, `CL_MEM_OBJECT_IMAGE2D_ARRAY` or `CL_MEM_OBJECT_IMAGE3D`.

`image_width` is the width of the image in pixels. For a 2D image and image array, the image width must be a value ≥ 1 and \leq `CL_DEVICE_IMAGE2D_MAX_WIDTH`. For a 3D image, the image width must be a value ≥ 1 and \leq `CL_DEVICE_IMAGE3D_MAX_WIDTH`. For a 1D image buffer, the image width must be a value ≥ 1 and \leq `CL_DEVICE_IMAGE_MAX_BUFFER_SIZE`. For a 1D image and 1D image array, the image width must be a value ≥ 1 and \leq `CL_DEVICE_IMAGE2D_MAX_WIDTH`.

`image_height` is height of the image in pixels. This is only used if the image is a 2D, 3D or 2D image array. For a 2D image or image array, the image height must be a value ≥ 1 and \leq `CL_DEVICE_IMAGE2D_MAX_HEIGHT`. For a 3D image, the image height must be a value ≥ 1 and \leq `CL_DEVICE_IMAGE3D_MAX_HEIGHT`.

`image_depth` is the depth of the image in pixels. This is only used if the image is a 3D image and must be a value ≥ 1 and \leq `CL_DEVICE_IMAGE3D_MAX_DEPTH`.

`image_array_size`¹¹ is the number of images in the image array. This is only used if the image is a 1D or 2D image array. The values for `image_array_size`, if specified, must be a value ≥ 1 and \leq `CL_DEVICE_IMAGE_MAX_ARRAY_SIZE`.

¹¹ Note that reading and writing 2D image arrays from a kernel with `image_array_size = 1` may be lower performance than 2D images.

`image_row_pitch` is the scan-line pitch in bytes. This must be 0 if `host_ptr` is NULL and can be either 0 or \geq `image_width` * size of element in bytes if `host_ptr` is not NULL. If `host_ptr` is not NULL and `image_row_pitch` = 0, `image_row_pitch` is calculated as `image_width` * size of element in bytes. If `image_row_pitch` is not 0, it must be a multiple of the image element size in bytes.

`image_slice_pitch` is the size in bytes of each 2D slice in the 3D image or the size in bytes of each image in a 1D or 2D image array. This must be 0 if `host_ptr` is NULL. If `host_ptr` is not NULL, `image_slice_pitch` can be either 0 or \geq `image_row_pitch` * `image_height` for a 2D image array or 3D image and can be either 0 or \geq `image_row_pitch` for a 1D image array. If `host_ptr` is not NULL and `image_slice_pitch` = 0, `image_slice_pitch` is calculated as `image_row_pitch` * `image_height` for a 2D image array or 3D image and `image_row_pitch` for a 1D image array. If `image_slice_pitch` is not 0, it must be a multiple of the `image_row_pitch`.

`num_mip_levels` and `num_samples` must be 0.

`buffer` refers to a valid buffer memory object if `image_type` is `CL_MEM_OBJECT_IMAGE1D_BUFFER`. Otherwise it must be NULL. For a 1D image buffer object, the image pixels are taken from the buffer object's data store. When the contents of a buffer object's data store are modified, those changes are reflected in the contents of the 1D image buffer object and vice-versa at corresponding synchronization points. The `image_width` * size of element in bytes must be \leq size of buffer object data store.

NOTE:

Concurrent reading from, writing to and copying between both a buffer object and 1D image buffer object associated with the buffer object is undefined. Only reading from both a buffer object and 1D image buffer object associated with the buffer object is defined.

5.3.2 Querying List of Supported Image Formats

The function

```
cl_int  clGetSupportedImageFormats (cl_context context,
                                     cl_mem_flags flags,
                                     cl_mem_object_type image_type,
                                     cl_uint num_entries,
                                     cl_image_format *image_formats,
                                     cl_uint *num_image_formats)
```

can be used to get the list of image formats supported by an OpenCL implementation when the following information about an image memory object is specified:

 Context

- ✚ Image type – 1D, 2D, or 3D image, 1D image buffer, 1D or 2D image array.
- ✚ Image object allocation information

clGetSupportedImageFormats returns a union of image formats supported by all devices in the context.

context is a valid OpenCL context on which the image object(s) will be created.

flags is a bit-field that is used to specify allocation and usage information about the image memory object being created and is described in *table 5.3*.

image_type describes the image type and must be either CL_MEM_OBJECT_IMAGE1D, CL_MEM_OBJECT_IMAGE1D_BUFFER, CL_MEM_OBJECT_IMAGE2D, CL_MEM_OBJECT_IMAGE3D, CL_MEM_OBJECT_IMAGE1D_ARRAY or CL_MEM_OBJECT_IMAGE2D_ARRAY.

num_entries specifies the number of entries that can be returned in the memory location given by *image_formats*.

image_formats is a pointer to a memory location where the list of supported image formats are returned. Each entry describes a *cl_image_format* structure supported by the OpenCL implementation. If *image_formats* is NULL, it is ignored.

num_image_formats is the actual number of supported image formats for a specific *context* and values specified by *flags*. If *num_image_formats* is NULL, it is ignored.

clGetSupportedImageFormats returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_CONTEXT if *context* is not a valid context.
- ✚ CL_INVALID_VALUE if *flags* or *image_type* are not valid, or if *num_entries* is 0 and *image_formats* is not NULL.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

If CL_DEVICE_IMAGE_SUPPORT specified in *table 4.3* is CL_TRUE, the values assigned to CL_DEVICE_MAX_READ_IMAGE_ARGS, CL_DEVICE_MAX_WRITE_IMAGE_ARGS, CL_DEVICE_IMAGE2D_MAX_WIDTH, CL_DEVICE_IMAGE2D_MAX_HEIGHT, CL_DEVICE_IMAGE3D_MAX_WIDTH, CL_DEVICE_IMAGE3D_MAX_HEIGHT, CL_DEVICE_IMAGE3D_MAX_DEPTH and CL_DEVICE_MAX_SAMPLERS by the implementation must be greater than or equal to the minimum values specified in *table 4.3*.

5.3.2.1 Minimum List of Supported Image Formats

For 1D, 2D, 3D image objects, 1D and 2D image array objects, the mandated minimum list of image formats that must be supported by all devices (for reading and writing) that support images is described in *table 5.8*.

| image_num_channels | image_channel_order | image_channel_data_type |
|---------------------------|----------------------------|--|
| 4 | CL_RGBA | CL_UNORM_INT8 CL_UNORM_INT16 CL_SIGNED_INT8 CL_SIGNED_INT16 CL_SIGNED_INT32 CL_UNSIGNED_INT8 CL_UNSIGNED_INT16 CL_UNSIGNED_INT32 CL_HALF_FLOAT CL_FLOAT |
| 4 | CL_BGRA | CL_UNORM_INT8 |

Table 5.8 *Min. list of supported image formats.*

5.3.3 Reading, Writing and Copying Image Objects

The following functions enqueue commands to read from an image or image array object to host memory or write to an image or image array object from host memory.

```
cl_int  clEnqueueReadImage (cl_command_queue command_queue,
                             cl_mem image,
                             cl_bool blocking_read,
                             const size_t *origin,
                             const size_t *region,
                             size_t row_pitch,
                             size_t slice_pitch,
                             void *ptr,
                             cl_uint num_events_in_wait_list,
                             const cl_event *event_wait_list,
                             cl_event *event)
```

```

cl_int  clEnqueueWriteImage (cl_command_queue command_queue,
                             cl_mem image,
                             cl_bool blocking_write,
                             const size_t *origin,
                             const size_t *region,
                             size_t input_row_pitch,
                             size_t input_slice_pitch,
                             const void * ptr,
                             cl_uint num_events_in_wait_list,
                             const cl_event *event_wait_list,
                             cl_event *event)

```

command_queue refers to the command-queue in which the read / write command will be queued. *command_queue* and *image* must be created with the same OpenCL context.

image refers to a valid image or image array object.

blocking_read and *blocking_write* indicate if the read and write operations are *blocking* or *non-blocking*.

If *blocking_read* is CL_TRUE i.e. the read command is blocking, **clEnqueueReadImage** does not return until the buffer data has been read and copied into memory pointed to by *ptr*.

If *blocking_read* is CL_FALSE i.e. the read command is non-blocking, **clEnqueueReadImage** queues a non-blocking read command and returns. The contents of the buffer that *ptr* points to cannot be used until the read command has completed. The *event* argument returns an event object which can be used to query the execution status of the read command. When the read command has completed, the contents of the buffer that *ptr* points to can be used by the application.

If *blocking_write* is CL_TRUE, the OpenCL implementation copies the data referred to by *ptr* and enqueues the write command in the command-queue. The memory pointed to by *ptr* can be reused by the application after the **clEnqueueWriteImage** call returns.

If *blocking_write* is CL_FALSE, the OpenCL implementation will use *ptr* to perform a non-blocking write. As the write is non-blocking the implementation can return immediately. The memory pointed to by *ptr* cannot be reused by the application after the call returns. The *event* argument returns an event object which can be used to query the execution status of the write command. When the write command has completed, the memory pointed to by *ptr* can then be reused by the application.

origin defines the (*x*, *y*, *z*) offset in pixels in the 1D, 2D or 3D image, the (*x*, *y*) offset and the image index in the 2D image array or the (*x*) offset and the image index in the 1D image array. If *image* is a 2D image object, *origin*[2] must be 0. If *image* is a 1D image or 1D image buffer object, *origin*[1] and *origin*[2] must be 0. If *image* is a 1D image array object, *origin*[2] must be

0. If *image* is a 1D image array object, *origin*[1] describes the image index in the 1D image array. If *image* is a 2D image array object, *origin*[2] describes the image index in the 2D image array.

region defines the (*width*, *height*, *depth*) in pixels of the 1D, 2D or 3D rectangle, the (*width*, *height*) in pixels of the 2D rectangle and the number of images of a 2D image array or the (*width*) in pixels of the 1D rectangle and the number of images of a 1D image array. If *image* is a 2D image object, *region*[2] must be 1. If *image* is a 1D image or 1D image buffer object, *region*[1] and *region*[2] must be 1. If *image* is a 1D image array object, *region*[2] must be 1. The values in *region* cannot be 0.

row_pitch in **clEnqueueReadImage** and *input_row_pitch* in **clEnqueueWriteImage** is the length of each row in bytes. This value must be greater than or equal to the element size in bytes * *width*. If *row_pitch* (or *input_row_pitch*) is set to 0, the appropriate row pitch is calculated based on the size of each element in bytes multiplied by *width*.

slice_pitch in **clEnqueueReadImage** and *input_slice_pitch* in **clEnqueueWriteImage** is the size in bytes of the 2D slice of the 3D region of a 3D image or each image of a 1D or 2D image array being read or written respectively. This must be 0 if *image* is a 1D or 2D image. This value must be greater than or equal to *row_pitch* * *height*. If *slice_pitch* (or *input_slice_pitch*) is set to 0, the appropriate slice pitch is calculated based on the *row_pitch* * *height*.

ptr is the pointer to a buffer in host memory where image data is to be read from or to be written to.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular read / write command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueReadImage and **clEnqueueWriteImage** return CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.

- ✦ CL_INVALID_CONTEXT if the context associated with *command_queue* and *image* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✦ CL_INVALID_MEM_OBJECT if *image* is not a valid image object.
- ✦ CL_INVALID_VALUE if the region being read or written specified by *origin* and *region* is out of bounds or if *ptr* is a NULL value.
- ✦ CL_INVALID_VALUE if values in *origin* and *region* do not follow rules described in the argument description for *origin* and *region*.
- ✦ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✦ CL_INVALID_IMAGE_SIZE if image dimensions (image width, height, specified or compute row and/or slice pitch) for *image* are not supported by device associated with *queue*.
- ✦ CL_IMAGE_FORMAT_NOT_SUPPORTED if image format (image channel order and data type) for *image* are not supported by device associated with *queue*.
- ✦ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with *image*.
- ✦ CL_INVALID_OPERATION if the device associated with *command_queue* does not support images (i.e. CL_DEVICE_IMAGE_SUPPORT specified in *table 4.3* is CL_FALSE).
- ✦ CL_INVALID_OPERATION if **clEnqueueReadImage** is called on *image* which has been created with CL_MEM_HOST_WRITE_ONLY or CL_MEM_HOST_NO_ACCESS.
- ✦ CL_INVALID_OPERATION if **clEnqueueWriteImage** is called on *image* which has been created with CL_MEM_HOST_READ_ONLY or CL_MEM_HOST_NO_ACCESS.
- ✦ CL_EXEC_STATUS_ERROR_FOR_EVENTS_IN_WAIT_LIST if the read and write operations are blocking and the execution status of any of the events in *event_wait_list* is a negative integer value.
- ✦ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✦ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

NOTE:

Calling **clEnqueueReadImage** to read a region of the *image* with the *ptr* argument value set to $host_ptr + (origin[2] * image\ slice\ pitch + origin[1] * image\ row\ pitch + origin[0] * bytes\ per\ pixel)$, where *host_ptr* is a pointer to the memory region specified when the *image* being read is created with CL_MEM_USE_HOST_PTR, must meet the following requirements in order to avoid undefined behavior:

- All commands that use this image object have finished execution before the read command begins execution.
- The *row_pitch* and *slice_pitch* argument values in **clEnqueueReadImage** must be set to the image row pitch and slice pitch.
- The image object is not mapped.
- The image object is not used by any command-queue until the read command has finished execution.

Calling **clEnqueueWriteImage** to update the latest bits in a region of the *image* with the *ptr* argument value set to $host_ptr + (origin[2] * image\ slice\ pitch + origin[1] * image\ row\ pitch + origin[0] * bytes\ per\ pixel)$, where *host_ptr* is a pointer to the memory region specified when the *image* being written is created with CL_MEM_USE_HOST_PTR, must meet the following requirements in order to avoid undefined behavior:

- The host memory region being written contains the latest bits when the enqueued write command begins execution.
- The *input_row_pitch* and *input_slice_pitch* argument values in **clEnqueueWriteImage** must be set to the image row pitch and slice pitch.
- The image object is not mapped.
- The image object is not used by any command-queue until the write command has finished execution.

The function

```
cl_int clEnqueueCopyImage (cl_command_queue command_queue,
                           cl_mem src_image,
                           cl_mem dst_image,
                           const size_t *src_origin,
                           const size_t *dst_origin,
                           const size_t *region,
                           cl_uint num_events_in_wait_list,
                           const cl_event *event_wait_list,
                           cl_event *event)
```

enqueues a command to copy image objects. *src_image* and *dst_image* can be 1D, 2D, 3D image or a 1D, 2D image array objects allowing us to perform the following actions:

- + Copy a 1D image object to a 1D image object.
- + Copy a 1D image object to a scanline of a 2D image object and vice-versa.
- + Copy a 1D image object to a scanline of a 2D slice of a 3D image object and vice-versa.
- + Copy a 1D image object to a scanline of a specific image index of a 1D or 2D image array object and vice-versa.
- + Copy a 2D image object to a 2D image object.
- + Copy a 2D image object to a 2D slice of a 3D image object and vice-versa.
- + Copy a 2D image object to a specific image index of a 2D image array object and vice-versa.
- + Copy images from a 1D image array object to a 1D image array object.
- + Copy images from a 2D image array object to a 2D image array object.
- + Copy a 3D image object to a 3D image object.

command_queue refers to the command-queue in which the copy command will be queued. The OpenCL context associated with *command_queue*, *src_image* and *dst_image* must be the same.

src_origin defines the (x, y, z) offset in pixels in the 1D, 2D or 3D image, the (x, y) offset and the image index in the 2D image array or the (x) offset and the image index in the 1D image array. If *image* is a 2D image object, *src_origin*[2] must be 0. If *src_image* is a 1D image object, *src_origin*[1] and *src_origin*[2] must be 0. If *src_image* is a 1D image array object, *src_origin*[2] must be 0. If *src_image* is a 1D image array object, *src_origin*[1] describes the image index in the 1D image array. If *src_image* is a 2D image array object, *src_origin*[2] describes the image index in the 2D image array.

dst_origin defines the (x, y, z) offset in pixels in the 1D, 2D or 3D image, the (x, y) offset and the image index in the 2D image array or the (x) offset and the image index in the 1D image array. If *dst_image* is a 2D image object, *dst_origin*[2] must be 0. If *dst_image* is a 1D image or 1D image buffer object, *dst_origin*[1] and *dst_origin*[2] must be 0. If *dst_image* is a 1D image array object, *dst_origin*[2] must be 0. If *dst_image* is a 1D image array object, *dst_origin*[1] describes the image index in the 1D image array. If *dst_image* is a 2D image array object, *dst_origin*[2] describes the image index in the 2D image array.

region defines the $(width, height, depth)$ in pixels of the 1D, 2D or 3D rectangle, the $(width, height)$ in pixels of the 2D rectangle and the number of images of a 2D image array or the $(width)$ in pixels of the 1D rectangle and the number of images of a 1D image array. If *src_image* or *dst_image* is a 2D image object, *region*[2] must be 1. If *src_image* or *dst_image* is a 1D image or 1D image buffer object, *region*[1] and *region*[2] must be 1. If *src_image* or *dst_image* is a 1D image array object, *region*[2] must be 1. The values in *region* cannot be 0.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list*

must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular copy command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. **clEnqueueBarrierWithWaitList** can be used instead. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

It is currently a requirement that the *src_image* and *dst_image* image memory objects for **clEnqueueCopyImage** must have the exact same image format (i.e. the *cl_image_format* descriptor specified when *src_image* and *dst_image* are created must match).

clEnqueueCopyImage returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_CONTEXT if the context associated with *command_queue*, *src_image* and *dst_image* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_MEM_OBJECT if *src_image* and *dst_image* are not valid image objects.
- ✚ CL_IMAGE_FORMAT_MISMATCH if *src_image* and *dst_image* do not use the same image format.
- ✚ CL_INVALID_VALUE if the 2D or 3D rectangular region specified by *src_origin* and *src_origin* + *region* refers to a region outside *src_image*, or if the 2D or 3D rectangular region specified by *dst_origin* and *dst_origin* + *region* refers to a region outside *dst_image*.
- ✚ CL_INVALID_VALUE if values in *src_origin*, *dst_origin* and *region* do not follow rules described in the argument description for *src_origin*, *dst_origin* and *region*.
- ✚ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✚ CL_INVALID_IMAGE_SIZE if image dimensions (image width, height, specified or compute row and/or slice pitch) for *src_image* or *dst_image* are not supported by device associated with *queue*.

- ✚ CL_IMAGE_FORMAT_NOT_SUPPORTED if image format (image channel order and data type) for *src_image* or *dst_image* are not supported by device associated with *queue*.
- ✚ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with *src_image* or *dst_image*.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.
- ✚ CL_INVALID_OPERATION if the device associated with *command_queue* does not support images (i.e. CL_DEVICE_IMAGE_SUPPORT specified in *table 4.3* is CL_FALSE).
- ✚ CL_MEM_COPY_OVERLAP if *src_image* and *dst_image* are the same image object and the source and destination regions overlap.

5.3.4 Filling Image Objects

The function

```
cl_int  clEnqueueFillImage (cl_command_queue command_queue,
                           cl_mem image,
                           const void *fill_color,
                           const size_t *origin,
                           const size_t *region,
                           cl_uint num_events_in_wait_list,
                           const cl_event *event_wait_list,
                           cl_event *event)
```

enqueues a command to fill an image object with a specified color. The usage information which indicates whether the memory object can be read or written by a kernel and/or the host and is given by the *cl_mem_flags* argument value specified when *image* is created is ignored by **clEnqueueFillImage**.

command_queue refers to the command-queue in which the fill command will be queued. The OpenCL context associated with *command_queue* and *image* must be the same.

image is a valid image object.

fill_color is the fill color. The fill color is a four component RGBA floating-point color value if the *image* channel data type is not an unnormalized signed and unsigned integer type, is a four component signed integer value if the *image* channel data type is an unnormalized signed integer type and is a four component unsigned integer value if the *image* channel data type is an

unnormalized unsigned integer type. The fill color will be converted to the appropriate image channel format and order associated with *image* as described in *sections 6.12.14* and *8.3*.

origin defines the (*x*, *y*, *z*) offset in pixels in the 1D, 2D or 3D image, the (*x*, *y*) offset and the image index in the 2D image array or the (*x*) offset and the image index in the 1D image array. If *image* is a 2D image object, *origin*[2] must be 0. If *image* is a 1D image or 1D image buffer object, *origin*[1] and *origin*[2] must be 0. If *image* is a 1D image array object, *origin*[2] must be 0. If *image* is a 1D image array object, *origin*[1] describes the image index in the 1D image array. If *image* is a 2D image array object, *origin*[2] describes the image index in the 2D image array.

region defines the (*width*, *height*, *depth*) in pixels of the 1D, 2D or 3D rectangle, the (*width*, *height*) in pixels of the 2D rectangle and the number of images of a 2D image array or the (*width*) in pixels of the 1D rectangle and the number of images of a 1D image array. If *image* is a 2D image object, *region*[2] must be 1. If *image* is a 1D image or 1D image buffer object, *region*[1] and *region*[2] must be 1. If *image* is a 1D image array object, *region*[2] must be 1. The values in *region* cannot be 0.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. **clEnqueueBarrierWithWaitList** can be used instead. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueFillImage returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_CONTEXT if the context associated with *command_queue* and *image* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_MEM_OBJECT if *image* is not a valid image object.
- ✚ CL_INVALID_VALUE if *fill_color* is NULL.

- ✦ CL_INVALID_VALUE if the region being filled as specified by *origin* and *region* is out of bounds or if *ptr* is a NULL value.
- ✦ CL_INVALID_VALUE if values in *origin* and *region* do not follow rules described in the argument description for *origin* and *region*.
- ✦ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✦ CL_INVALID_IMAGE_SIZE if image dimensions (image width, height, specified or compute row and/or slice pitch) for *image* are not supported by device associated with *queue*.
- ✦ CL_IMAGE_FORMAT_NOT_SUPPORTED if image format (image channel order and data type) for *image* are not supported by device associated with *queue*.
- ✦ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with *image*.
- ✦ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✦ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.3.5 Copying between Image and Buffer Objects

The function

```
cl_int  clEnqueueCopyImageToBuffer (cl_command_queue command_queue,
                                     cl_mem src_image,
                                     cl_mem dst_buffer,
                                     const size_t *src_origin,
                                     const size_t *region,
                                     size_t dst_offset,
                                     cl_uint num_events_in_wait_list,
                                     const cl_event *event_wait_list,
                                     cl_event *event)
```

enqueues a command to copy an image object to a buffer object.

command_queue must be a valid command-queue. The OpenCL context associated with *command_queue*, *src_image* and *dst_buffer* must be the same.

src_image is a valid image object.

dst_buffer is a valid buffer object.

src_origin defines the (*x*, *y*, *z*) offset in pixels in the 1D, 2D or 3D image, the (*x*, *y*) offset and the image index in the 2D image array or the (*x*) offset and the image index in the 1D image array. If *src_image* is a 2D image object, *src_origin*[2] must be 0. If *src_image* is a 1D image or 1D image buffer object, *src_origin*[1] and *src_origin*[2] must be 0. If *src_image* is a 1D image array object, *src_origin*[2] must be 0. If *src_image* is a 1D image array object, *src_origin*[1] describes the image index in the 1D image array. If *src_image* is a 2D image array object, *src_origin*[2] describes the image index in the 2D image array.

region defines the (*width*, *height*, *depth*) in pixels of the 1D, 2D or 3D rectangle, the (*width*, *height*) in pixels of the 2D rectangle and the number of images of a 2D image array or the (*width*) in pixels of the 1D rectangle and the number of images of a 1D image array. If *src_image* is a 2D image object, *region*[2] must be 1. If *src_image* is a 1D image or 1D image buffer object, *region*[1] and *region*[2] must be 1. If *src_image* is a 1D image array object, *region*[2] must be 1. The values in *region* cannot be 0.

dst_offset refers to the offset where to begin copying data into *dst_buffer*. The size in bytes of the region to be copied referred to as *dst_cb* is computed as *width* * *height* * *depth* * *bytes/image element* if *src_image* is a 3D image object, is computed as *width* * *height* * *bytes/image element* if *src_image* is a 2D image, is computed as *width* * *height* * *arraysize* * *bytes/image element* if *src_image* is a 2D image array object, is computed as *width* * *bytes/image element* if *src_image* is a 1D image or 1D image buffer object and is computed as *width* * *arraysize* * *bytes/image element* if *src_image* is a 1D image array object.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular copy command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. **clEnqueueBarrierWithWaitList** can be used instead. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueCopyImageToBuffer returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.

- ✦ CL_INVALID_CONTEXT if the context associated with *command_queue*, *src_image* and *dst_buffer* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✦ CL_INVALID_MEM_OBJECT if *src_image* is not a valid image object or *dst_buffer* is not a valid buffer object or if *src_image* is a 1D image buffer object created from *dst_buffer*.
- ✦ CL_INVALID_VALUE if the 1D, 2D or 3D rectangular region specified by *src_origin* and *src_origin + region* refers to a region outside *src_image*, or if the region specified by *dst_offset* and *dst_offset + dst_cb* to a region outside *dst_buffer*.
- ✦ CL_INVALID_VALUE if values in *src_origin* and *region* do not follow rules described in the argument description for *src_origin* and *region*.
- ✦ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✦ CL_MISALIGNED_SUB_BUFFER_OFFSET if *dst_buffer* is a sub-buffer object and *offset* specified when the sub-buffer object is created is not aligned to CL_DEVICE_MEM_BASE_ADDR_ALIGN value for device associated with *queue*.
- ✦ CL_INVALID_IMAGE_SIZE if image dimensions (image width, height, specified or compute row and/or slice pitch) for *src_image* are not supported by device associated with *queue*.
- ✦ CL_IMAGE_FORMAT_NOT_SUPPORTED if image format (image channel order and data type) for *src_image* are not supported by device associated with *queue*.
- ✦ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with *src_image* or *dst_buffer*.
- ✦ CL_INVALID_OPERATION if the device associated with *command_queue* does not support images (i.e. CL_DEVICE_IMAGE_SUPPORT specified in *table 4.3* is CL_FALSE).
- ✦ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✦ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```
cl_int  clEnqueueCopyBufferToImage (cl_command_queue command_queue,
                                     cl_mem src_buffer,
                                     cl_mem dst_image,
                                     size_t src_offset,
                                     const size_t *dst_origin,
                                     const size_t *region,
                                     cl_uint num_events_in_wait_list,
                                     const cl_event *event_wait_list,
                                     cl_event *event)
```

enqueues a command to copy a buffer object to an image object.

command_queue must be a valid command-queue. The OpenCL context associated with *command_queue*, *src_buffer* and *dst_image* must be the same.

src_buffer is a valid buffer object.

dst_image is a valid image object.

src_offset refers to the offset where to begin copying data from *src_buffer*.

dst_origin defines the (*x*, *y*, *z*) offset in pixels in the 1D, 2D or 3D image, the (*x*, *y*) offset and the image index in the 2D image array or the (*x*) offset and the image index in the 1D image array. If *dst_image* is a 2D image object, *dst_origin*[2] must be 0. If *dst_image* is a 1D image or 1D image buffer object, *dst_origin*[1] and *dst_origin*[2] must be 0. If *dst_image* is a 1D image array object, *dst_origin*[2] must be 0. If *dst_image* is a 1D image array object, *dst_origin*[1] describes the image index in the 1D image array. If *dst_image* is a 2D image array object, *dst_origin*[2] describes the image index in the 2D image array.

region defines the (*width*, *height*, *depth*) in pixels of the 1D, 2D or 3D rectangle, the (*width*, *height*) in pixels of the 2D rectangle and the number of images of a 2D image array or the (*width*) in pixels of the 1D rectangle and the number of images of a 1D image array. If *dst_image* is a 2D image object, *region*[2] must be 1. If *dst_image* is a 1D image or 1D image buffer object, *region*[1] and *region*[2] must be 1. If *dst_image* is a 1D image array object, *region*[2] must be 1. The values in *region* cannot be 0.

The size in bytes of the region to be copied from *src_buffer* referred to as *src_cb* is computed as *width* * *height* * *depth* * *bytes/image element* if *dst_image* is a 3D image object, is computed as *width* * *height* * *bytes/image element* if *dst_image* is a 2D image, is computed as *width* * *height* * *arraysize* * *bytes/image element* if *dst_image* is a 2D image array object, is computed as *width* * *bytes/image element* if *dst_image* is a 1D image or 1D image buffer object and is computed as *width* * *arraysize* * *bytes/image element* if *dst_image* is a 1D image array object.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular copy command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. **clEnqueueBarrierWithWaitList** can be used instead. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueCopyBufferToImage returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_CONTEXT if the context associated with *command_queue*, *src_buffer* and *dst_image* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_MEM_OBJECT if *src_buffer* is not a valid buffer object or *dst_image* is not a valid image object or if *dst_image* is a 1D image buffer object created from *src_buffer*.
- ✚ CL_INVALID_VALUE if the 1D, 2D or 3D rectangular region specified by *dst_origin* and *dst_origin + region* refer to a region outside *dst_image*, or if the region specified by *src_offset* and *src_offset + src_cb* refer to a region outside *src_buffer*.
- ✚ CL_INVALID_VALUE if values in *dst_origin* and *region* do not follow rules described in the argument description for *dst_origin* and *region*.
- ✚ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✚ CL_MISALIGNED_SUB_BUFFER_OFFSET if *src_buffer* is a sub-buffer object and *offset* specified when the sub-buffer object is created is not aligned to CL_DEVICE_MEM_BASE_ADDR_ALIGN value for device associated with *queue*.
- ✚ CL_INVALID_IMAGE_SIZE if image dimensions (image width, height, specified or compute row and/or slice pitch) for *dst_image* are not supported by device associated with *queue*.

- ✚ CL_IMAGE_FORMAT_NOT_SUPPORTED if image format (image channel order and data type) for *dst_image* are not supported by device associated with *queue*.
- ✚ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with *src_buffer* or *dst_image*.
- ✚ CL_INVALID_OPERATION if the device associated with *command_queue* does not support images (i.e. CL_DEVICE_IMAGE_SUPPORT specified in *table 4.3* is CL_FALSE).
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.3.6 Mapping Image Objects

The function

```
void * clEnqueueMapImage (cl_command_queue command_queue,
                          cl_mem image,
                          cl_bool blocking_map,
                          cl_map_flags map_flags,
                          const size_t *origin,
                          const size_t *region,
                          size_t *image_row_pitch,
                          size_t *image_slice_pitch,
                          cl_uint num_events_in_wait_list,
                          const cl_event *event_wait_list,
                          cl_event *event,
                          cl_int *errcode_ret)
```

enqueues a command to map a region in the image object given by *image* into the host address space and returns a pointer to this mapped region.

command_queue must be a valid command-queue.

image is a valid image object. The OpenCL context associated with *command_queue* and *image* must be the same.

blocking_map indicates if the map operation is *blocking* or *non-blocking*.

If *blocking_map* is CL_TRUE, **clEnqueueMapImage** does not return until the specified region in *image* is mapped into the host address space and the application can access the contents of the mapped region using the pointer returned by **clEnqueueMapImage**.

If *blocking_map* is CL_FALSE i.e. map operation is non-blocking, the pointer to the mapped region returned by **clEnqueueMapImage** cannot be used until the map command has completed. The *event* argument returns an event object which can be used to query the execution status of the map command. When the map command is completed, the application can access the contents of the mapped region using the pointer returned by **clEnqueueMapImage**.

map_flags is a bit-field and is described in *table 5.5*.

origin defines the (*x*, *y*, *z*) offset in pixels in the 1D, 2D or 3D image, the (*x*, *y*) offset and the image index in the 2D image array or the (*x*) offset and the image index in the 1D image array. If *image* is a 2D image object, *origin*[2] must be 0. If *image* is a 1D image or 1D image buffer object, *origin*[1] and *origin*[2] must be 0. If *image* is a 1D image array object, *origin*[2] must be 0. If *image* is a 1D image array object, *origin*[1] describes the image index in the 1D image array. If *image* is a 2D image array object, *origin*[2] describes the image index in the 2D image array.

region defines the (*width*, *height*, *depth*) in pixels of the 1D, 2D or 3D rectangle, the (*width*, *height*) in pixels of the 2D rectangle and the number of images of a 2D image array or the (*width*) in pixels of the 1D rectangle and the number of images of a 1D image array. If *image* is a 2D image object, *region*[2] must be 1. If *image* is a 1D image or 1D image buffer object, *region*[1] and *region*[2] must be 1. If *image* is a 1D image array object, *region*[2] must be 1. The values in *region* cannot be 0.

image_row_pitch returns the scan-line pitch in bytes for the mapped region. This must be a non-NULL value.

image_slice_pitch returns the size in bytes of each 2D slice of a 3D image or the size of each 1D or 2D image in a 1D or 2D image array for the mapped region. For a 1D and 2D image, zero is returned if this argument is not NULL. For a 3D image, 1D and 2D image array, *image_slice_pitch* must be a non-NULL value.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before **clEnqueueMapImage** can be executed. If *event_wait_list* is NULL, then **clEnqueueMapImage** does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

errcode_ret will return an appropriate error code. If *errcode_ret* is NULL, no error code is returned.

clEnqueueMapImage will return a pointer to the mapped region. The *errcode_ret* is set to CL_SUCCESS.

A NULL pointer is returned otherwise with one of the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_CONTEXT if context associated with *command_queue* and *image* are not the same or if context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_MEM_OBJECT if *image* is not a valid image object.
- ✚ CL_INVALID_VALUE if region being mapped given by (*origin*, *origin+region*) is out of bounds or if values specified in *map_flags* are not valid.
- ✚ CL_INVALID_VALUE if values in *origin* and *region* do not follow rules described in the argument description for *origin* and *region*.
- ✚ CL_INVALID_VALUE if *image_row_pitch* is NULL.
- ✚ CL_INVALID_VALUE if *image* is a 3D image, 1D or 2D image array object and *image_slice_pitch* is NULL.
- ✚ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✚ CL_INVALID_IMAGE_SIZE if image dimensions (image width, height, specified or compute row and/or slice pitch) for *image* are not supported by device associated with *queue*.
- ✚ CL_IMAGE_FORMAT_NOT_SUPPORTED if image format (image channel order and data type) for *image* are not supported by device associated with *queue*.
- ✚ CL_MAP_FAILURE if there is a failure to map the requested region into the host address space. This error cannot occur for image objects created with CL_MEM_USE_HOST_PTR or CL_MEM_ALLOC_HOST_PTR.
- ✚ CL_EXEC_STATUS_ERROR_FOR_EVENTS_IN_WAIT_LIST if the map operation is blocking and the execution status of any of the events in *event_wait_list* is a negative integer value.

- ✦ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with *image*.
- ✦ CL_INVALID_OPERATION if the device associated with *command_queue* does not support images (i.e. CL_DEVICE_IMAGE_SUPPORT specified in *table 4.3* is CL_FALSE).
- ✦ CL_INVALID_OPERATION if *image* has been created with CL_MEM_HOST_WRITE_ONLY or CL_MEM_HOST_NO_ACCESS and CL_MAP_READ is set in *map_flags* or if *image* has been created with CL_MEM_HOST_READ_ONLY or CL_MEM_HOST_NO_ACCESS and CL_MAP_WRITE or CL_MAP_WRITE_INVALIDATE_REGION is set in *map_flags*.
- ✦ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✦ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The pointer returned maps a 1D, 2D or 3D region starting at *origin* and is at least *region[0]* pixels in size for a 1D image, 1D image buffer or 1D image array, (*image_row_pitch* * *region[1]*) pixels in size for a 2D image or 2D image array, and (*image_slice_pitch* * *region[2]*) pixels in size for a 3D image. The result of a memory access outside this region is undefined.

If the image object is created with CL_MEM_USE_HOST_PTR set in *mem_flags*, the following will be true:

- ✦ The *host_ptr* specified in **clCreateImage** is guaranteed to contain the latest bits in the region being mapped when the **clEnqueueMapImage** command has completed.
- ✦ The pointer value returned by **clEnqueueMapImage** will be derived from the *host_ptr* specified when the image object is created.

Mapped image objects are unmapped using **clEnqueueUnmapMemObject**. This is described in *section 5.4.2*.

5.3.7 Image Object Queries

To get information that is common to all memory objects (buffer and image objects), use the **clGetMemObjectInfo** function described in *section 5.4.5*.

To get information specific to an image object created with **clCreateImage**, use the following function

```

cl_int      clGetImageInfo (cl_mem image,
                             cl_image_info param_name,
                             size_t param_value_size,
                             void *param_value,
                             size_t *param_value_size_ret)

```

image specifies the image object being queried.

param_name specifies the information to query. The list of supported *param_name* types and the information returned in *param_value* by **clGetImageInfo** is described in *table 5.9*.

param_value is a pointer to memory where the appropriate result being queried is returned. If *param_value* is NULL, it is ignored.

param_value_size is used to specify the size in bytes of memory pointed to by *param_value*. This size must be \geq size of return type as described in *table 5.9*.

param_value_size_ret returns the actual size in bytes of data being queried by *param_value*. If *param_value_size_ret* is NULL, it is ignored.

clGetImageInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_VALUE if *param_name* is not valid, or if size in bytes specified by *param_value_size* is $<$ size of return type as described in *table 5.9* and *param_value* is not NULL.
- ✚ CL_INVALID_MEM_OBJECT if *image* is a not a valid image object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

| cl_image_info | Return type | Info. returned in <i>param_value</i> |
|-----------------------|-----------------|---|
| CL_IMAGE_FORMAT | cl_image_format | Return image format descriptor specified when <i>image</i> is created with clCreateImage . |
| CL_IMAGE_ELEMENT_SIZE | size_t | Return size of each element of the image memory object given by <i>image</i> . An element is made up of <i>n</i> channels. The value of <i>n</i> is given in <i>cl_image_format</i> descriptor. |

| | | |
|--------------------------------|---------|--|
| CL_IMAGE_ROW_PITCH | size_t | Return calculated row pitch in bytes of a row of elements of the image object given by <i>image</i> . |
| CL_IMAGE_SLICE_PITCH | size_t | Return calculated slice pitch in bytes of a 2D slice for the 3D image object or size of each image in a 1D or 2D image array given by <i>image</i> . For a 1D image, 1D image buffer and 2D image object return 0. |
| CL_IMAGE_WIDTH | size_t | Return width of the image in pixels. |
| CL_IMAGE_HEIGHT | size_t | Return height of the image in pixels. For a 1D image, 1D image buffer and 1D image array object, height = 0. |
| CL_IMAGE_DEPTH | size_t | Return depth of the image in pixels. For a 1D image, 1D image buffer, 2D image or 1D and 2D image array object, depth = 0. |
| CL_IMAGE_ARRAY_SIZE | size_t | Return number of images in the image array. If <i>image</i> is not an image array, 0 is returned. |
| CL_IMAGE_BUFFER | cl_mem | Return buffer object associated with <i>image</i> . |
| CL_IMAGE_NUM_MIP_LEVELS | cl_uint | Return num_mip_levels associated with <i>image</i> . |
| CL_IMAGE_NUM_SAMPLES | cl_uint | Return num_samples associated with <i>image</i> . |

Table 5.9 List of supported param_names by clGetImageInfo

5.4 Querying, Unmapping, Migrating, Retaining and Releasing Memory Objects

5.4.1 Retaining and Releasing Memory Objects

The function

```
cl_int          clRetainMemObject (cl_mem memobj)
```

increments the *memobj* reference count. **clRetainMemObject** returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_MEM_OBJECT if *memobj* is not a valid memory object (buffer or image object).
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

clCreateBuffer, **clCreateSubBuffer** and **clCreateImage** perform an implicit retain.

The function

```
cl_int          clReleaseMemObject (cl_mem memobj)
```

decrements the *memobj* reference count. **clReleaseMemObject** returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_MEM_OBJECT if *memobj* is not a valid memory object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

After the *memobj* reference count becomes zero and commands queued for execution on a command-queue(s) that use *memobj* have finished, the memory object is deleted. If *memobj* is a buffer object, *memobj* cannot be deleted until all sub-buffer objects associated with *memobj* are deleted.

The function

```
cl_int          clSetMemObjectDestructorCallback (cl_mem memobj,
                                                void (CL_CALLBACK *pfn_notify)(cl_mem memobj,
                                                                    void *user_data),
                                                void *user_data)
```

registers a user callback function with a memory object. Each call to **clSetMemObjectDestructorCallback** registers the specified user callback function on a callback stack associated with *memobj*. The registered user callback functions are called in the reverse order in which they were registered. The user callback functions are called and then the memory object's resources are freed and the memory object is deleted. This provides a mechanism for the application (and libraries) using *memobj* to be notified when the memory referenced by *host_ptr*, specified when the memory object is created and used as the storage bits for the memory object, can be reused or freed.

memobj is a valid memory object.

pfn_notify is the callback function that can be registered by the application. This callback function may be called asynchronously by the OpenCL implementation. It is the application's responsibility to ensure that the callback function is thread-safe. The parameters to this callback function are:

- ✚ *memobj* is the memory object being deleted. When the user callback is called by the implementation, this memory object is no longer valid. *memobj* is only provided for reference purposes.
- ✚ *user_data* is a pointer to user supplied data.

user_data will be passed as the *user_data* argument when *pfn_notify* is called. *user_data* can be NULL.

clSetMemObjectDestructorCallback returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_MEM_OBJECT if *memobj* is not a valid memory object.
- ✚ CL_INVALID_VALUE if *pfn_notify* is NULL.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

NOTE: When the user callback function is called by the implementation, the contents of the memory region pointed to by *host_ptr* (if the memory object is created with CL_MEM_USE_HOST_PTR) are undefined. The callback function is typically used by the application to either free or reuse the memory region pointed to by *host_ptr*.

The behavior of calling expensive system routines, OpenCL API calls to create contexts or command-queues, or blocking OpenCL operations from the following list below, in a callback is undefined.

clFinish,
clWaitForEvents,
blocking calls to **clEnqueueReadBuffer**, **clEnqueueReadBufferRect**,
clEnqueueWriteBuffer, **clEnqueueWriteBufferRect**,
blocking calls to **clEnqueueReadImage** and **clEnqueueWriteImage**,
blocking calls to **clEnqueueMapBuffer**,
clEnqueueMapImage,
blocking calls to **clBuildProgram**, **clCompileProgram** or **clLinkProgram**

If an application needs to wait for completion of a routine from the above list in a callback, please use the non-blocking form of the function, and assign a completion callback to it to do the remainder of your work. Note that when a callback (or other code) enqueues commands to a command-queue, the commands are not required to begin execution until the queue is flushed. In standard usage, blocking enqueue calls serve this role by implicitly flushing the queue. Since blocking calls are not permitted in callbacks, those callbacks that enqueue commands on a command queue should either call **clFlush** on the queue before returning or arrange for **clFlush** to be called later on another thread.

The user callback function may not call OpenCL APIs with the memory object for which the callback function is invoked and for such cases the behavior of OpenCL APIs is considered to be undefined.

5.4.2 Unmapping Mapped Memory Objects

The function

```
cl_int clEnqueueUnmapMemObject (cl_command_queue command_queue,
                                cl_mem memobj,
                                void *mapped_ptr,
                                cl_uint num_events_in_wait_list,
                                const cl_event *event_wait_list,
                                cl_event *event)
```

enqueues a command to unmap a previously mapped region of a memory object. Reads or writes from the host using the pointer returned by **clEnqueueMapBuffer** or **clEnqueueMapImage** are considered to be complete.

command_queue must be a valid command-queue.

memobj is a valid memory object. The OpenCL context associated with *command_queue* and *memobj* must be the same.

mapped_ptr is the host address returned by a previous call to **clEnqueueMapBuffer**, or **clEnqueueMapImage** for *memobj*.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before **clEnqueueUnmapMemObject** can be executed. If *event_wait_list* is NULL, then **clEnqueueUnmapMemObject** does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. **clEnqueueBarrierWithWaitList** can be used instead. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueUnmapMemObject returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_MEM_OBJECT if *memobj* is not a valid memory object.
- ✚ CL_INVALID_VALUE if *mapped_ptr* is not a valid pointer returned by **clEnqueueMapBuffer**, or **clEnqueueMapImage** for *memobj*.
- ✚ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or if *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

- ✚ CL_INVALID_CONTEXT if context associated with *command_queue* and *memobj* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.

clEnqueueMapBuffer, and **clEnqueueMapImage** increments the mapped count of the memory object. The initial mapped count value of the memory object is zero. Multiple calls to **clEnqueueMapBuffer**, or **clEnqueueMapImage** on the same memory object will increment this mapped count by appropriate number of calls. **clEnqueueUnmapMemObject** decrements the mapped count of the memory object.

clEnqueueMapBuffer, and **clEnqueueMapImage** act as synchronization points for a region of the buffer object being mapped.

5.4.3 Accessing mapped regions of a memory object

This section describes the behavior of OpenCL commands that access mapped regions of a memory object.

The contents of the region of a memory object and associated memory objects (sub-buffer objects or 1D image buffer objects that overlap this region) mapped for writing (i.e. CL_MAP_WRITE or CL_MAP_WRITE_INVALIDATE_REGION is set in *map_flags* argument to **clEnqueueMapBuffer**, or **clEnqueueMapImage**) are considered to be undefined until this region is unmapped.

Multiple commands in command-queues can map a region or overlapping regions of a memory object and associated memory objects (sub-buffer objects or 1D image buffer objects that overlap this region) for reading (i.e. *map_flags* = CL_MAP_READ). The contents of the regions of a memory object mapped for reading can also be read by kernels and other OpenCL commands (such as **clEnqueueCopyBuffer**) executing on a device(s).

Mapping (and unmapping) overlapped regions in a memory object and/or associated memory objects (sub-buffer objects or 1D image buffer objects that overlap this region) for writing is an error and will result in CL_INVALID_OPERATION error returned by **clEnqueueMapBuffer**, or **clEnqueueMapImage**.

If a memory object is currently mapped for writing, the application must ensure that the memory object is unmapped before any enqueued kernels or commands that read from or write to this memory object or any of its associated memory objects (sub-buffer or 1D image buffer objects) or its parent object (if the memory object is a sub-buffer or 1D image buffer object) begin execution; otherwise the behavior is undefined.

If a memory object is currently mapped for reading, the application must ensure that the memory object is unmapped before any enqueued kernels or commands that write to this memory object or any of its associated memory objects (sub-buffer or 1D image buffer objects) or its parent object (if the memory object is a sub-buffer or 1D image buffer object) begin execution;

otherwise the behavior is undefined.

Accessing the contents of the memory region referred to by the mapped pointer that has been unmapped is undefined.

The mapped pointer returned by **clEnqueueMapBuffer** or **clEnqueueMapImage** can be used as *ptr* argument value to **clEnqueue{Read | Write}Buffer**, **clEnqueue{Read | Write}BufferRect**, **clEnqueue{Read | Write}Image** provided the rules described above are adhered to.

5.4.4 Migrating Memory Objects

This section describes a mechanism for assigning which device an OpenCL memory object resides. A user may wish to have more explicit control over the location of their memory objects on creation. This could be used to:

- ✚ Ensure that an object is allocated on a specific device prior to usage.
- ✚ Preemptively migrate an object from one device to another.

The function

```
cl_int clEnqueueMigrateMemObjects (cl_command_queue command_queue,
                                   cl_uint num_mem_objects,
                                   const cl_mem *mem_objects,
                                   cl_mem_migration_flags flags,
                                   cl_uint num_events_in_wait_list,
                                   const cl_event *event_wait_list,
                                   cl_event *event)
```

enqueues a command to indicate which device a set of memory objects should be associated with. Typically, memory objects are implicitly migrated to a device for which enqueued commands, using the memory object, are targeted. **clEnqueueMigrateMemObjects** allows this migration to be explicitly performed ahead of the dependent commands. This allows a user to preemptively change the association of a memory object, through regular command queue scheduling, in order to prepare for another upcoming command. This also permits an application to overlap the placement of memory objects with other unrelated operations before these memory objects are needed potentially hiding transfer latencies. Once the event, returned from **clEnqueueMigrateMemObjects**, has been marked CL_COMPLETE the memory objects specified in *mem_objects* have been successfully migrated to the device associated with *command_queue*. The migrated memory object shall remain resident on the device until another command is enqueued that either implicitly or explicitly migrates it away.

clEnqueueMigrateMemObjects can also be used to direct the initial placement of a memory object, after creation, possibly avoiding the initial overhead of instantiating the object on the first enqueued command to use it.

The user is responsible for managing the event dependencies, associated with this command, in order to avoid overlapping access to memory objects. Improperly specified event dependencies passed to **clEnqueueMigrateMemObjects** could result in undefined results.

command_queue is a valid command-queue. The specified set of memory objects in *mem_objects* will be migrated to the OpenCL device associated with *command_queue* or to the host if the CL_MIGRATE_MEM_OBJECT_HOST has been specified.

num_mem_objects is the number of memory objects specified in *mem_objects*.

mem_objects is a pointer to a list of memory objects.

flags is a bit-field that is used to specify migration options. The following table describes the possible values for flags.

| cl_mem_migration flags | Description |
|--|---|
| CL_MIGRATE_MEM_OBJECT_HOST | This flag indicates that the specified set of memory objects are to be migrated to the host, regardless of the target command-queue. |
| CL_MIGRATE_MEM_OBJECT_CONTENT_UNDEFINED | This flag indicates that the contents of the set of memory objects are undefined after migration. The specified set of memory objects are migrated to the device associated with <i>command_queue</i> without incurring the overhead of migrating their contents. |

Table 5.10 Supported *cl_mem_migration flags*.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular command and can be used to query or queue a wait for this particular command to complete. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueMigrateMemObjects return CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✦ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✦ CL_INVALID_CONTEXT if the context associated with *command_queue* and memory objects in *mem_objects* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✦ CL_INVALID_MEM_OBJECT if any of the memory objects in *mem_objects* is not a valid memory object.
- ✦ CL_INVALID_VALUE if *num_mem_objects* is zero or if *mem_objects* is NULL.
- ✦ CL_INVALID_VALUE if *flags* is not 0 or is not any of the values described in the table above.
- ✦ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✦ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for the specified set of memory objects in *mem_objects*.
- ✦ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✦ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.4.5 Memory Object Queries

To get information that is common to all memory objects (buffer and image objects), use the following function

```
cl_int          clGetMemObjectInfo (cl_mem memobj,
                                     cl_mem_info param_name,
                                     size_t param_value_size,
                                     void *param_value,
                                     size_t *param_value_size_ret)
```

memobj specifies the memory object being queried.

param_name specifies the information to query. The list of supported *param_name* types and the information returned in *param_value* by **clGetMemObjectInfo** is described in *table 5.11*.

param_value is a pointer to memory where the appropriate result being queried is returned. If *param_value* is NULL, it is ignored.

param_value_size is used to specify the size in bytes of memory pointed to by *param_value*. This size must be \geq size of return type as described in *table 5.11*.

param_value_size_ret returns the actual size in bytes of data being queried by *param_value*. If *param_value_size_ret* is NULL, it is ignored.

clGetMemObjectInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_VALUE if *param_name* is not valid, or if size in bytes specified by *param_value_size* is $<$ size of return type as described in *table 5.11* and *param_value* is not NULL.
- ✚ CL_INVALID_MEM_OBJECT if *memobj* is a not a valid memory object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

| cl_mem_info | Return type | Info. returned in <i>param_value</i> |
|---------------------|--------------------|--|
| CL_MEM_TYPE | cl_mem_object_type | Returns one of the following values: CL_MEM_OBJECT_BUFFER if <i>memobj</i> is created with clCreateBuffer or clCreateSubBuffer . cl_image_desc.image_type argument value if <i>memobj</i> is created with clCreateImage . |
| CL_MEM_FLAGS | cl_mem_flags | Return the <i>flags</i> argument value specified when <i>memobj</i> is created with clCreateBuffer , clCreateSubBuffer or clCreateImage . If <i>memobj</i> is a sub-buffer the memory access qualifiers inherited from parent buffer is also returned. |
| CL_MEM_SIZE | size_t | Return actual size of the data store associated with <i>memobj</i> in bytes. |

| | | |
|--|------------|---|
| CL_MEM_HOST_PTR | void * | <p>If <i>memobj</i> is created with clCreateBuffer or clCreateImage and CL_MEM_USE_HOST_PTR is specified in <i>mem_flags</i>, return the <i>host_ptr</i> argument value specified when <i>memobj</i> is created. Otherwise a NULL value is returned.</p> <p>If <i>memobj</i> is created with clCreateSubBuffer, return the <i>host_ptr</i> + <i>origin</i> value specified when <i>memobj</i> is created. <i>host_ptr</i> is the argument value specified to clCreateBuffer and CL_MEM_USE_HOST_PTR is specified in <i>mem_flags</i> for memory object from which <i>memobj</i> is created. Otherwise a NULL value is returned.</p> |
| CL_MEM_MAP_COUNT¹² | cl_uint | Map count. |
| CL_MEM_REFERENCE_COUNT¹³ | cl_uint | Return <i>memobj</i> reference count. |
| CL_MEM_CONTEXT | cl_context | Return context specified when memory object is created. If <i>memobj</i> is created using clCreateSubBuffer , the context associated with the memory object specified as the <i>buffer</i> argument to clCreateSubBuffer is returned. |
| CL_MEM_ASSOCIATED_MEMOBJECT | cl_mem | <p>Return memory object from which <i>memobj</i> is created.</p> <p>This returns the memory object specified as <i>buffer</i> argument to clCreateSubBuffer.</p> <p>Otherwise a NULL value is returned.</p> |
| CL_MEM_OFFSET | size_t | <p>Return offset if <i>memobj</i> is a sub-buffer object created using clCreateSubBuffer.</p> <p>This return 0 if <i>memobj</i> is not a sub-buffer object.</p> |

Table 5.11 List of supported *param_names* by *clGetMemObjectInfo*

¹² The map count returned should be considered immediately stale. It is unsuitable for general use in applications. This feature is provided for debugging.

¹³ The reference count returned should be considered immediately stale. It is unsuitable for general use in applications. This feature is provided for identifying memory leaks.

5.5 Sampler Objects

A sampler object describes how to sample an image when the image is read in the kernel. The built-in functions to read from an image in a kernel take a sampler as an argument. The sampler arguments to the image read function can be sampler objects created using OpenCL functions and passed as argument values to the kernel or can be samplers declared inside a kernel. In this section we discuss how sampler objects are created using OpenCL functions.

5.5.1 Creating Sampler Objects

The function

```
cl_sampler    clCreateSampler (cl_context context,
                              cl_bool normalized_coords,
                              cl_addressing_mode addressing_mode,
                              cl_filter_mode filter_mode,
                              cl_int *errcode_ret)
```

creates a sampler object. Refer to *section 6.12.14.1* for a detailed description of how samplers work.

context must be a valid OpenCL context.

normalized_coords determines if the image coordinates specified are normalized (if *normalized_coords* is CL_TRUE) or not (if *normalized_coords* is CL_FALSE).

addressing_mode specifies how out-of-range image coordinates are handled when reading from an image. This can be set to CL_ADDRESS_MIRRORED_REPEAT, CL_ADDRESS_REPEAT, CL_ADDRESS_CLAMP_TO_EDGE, CL_ADDRESS_CLAMP and CL_ADDRESS_NONE.

filter_mode specifies the type of filter that must be applied when reading an image. This can be CL_FILTER_NEAREST, or CL_FILTER_LINEAR.

errcode_ret will return an appropriate error code. If *errcode_ret* is NULL, no error code is returned.

clCreateSampler returns a valid non-zero sampler object and *errcode_ret* is set to CL_SUCCESS if the sampler object is created successfully. Otherwise, it returns a NULL value with one of the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_CONTEXT if *context* is not a valid context.
- ✚ CL_INVALID_VALUE if *addressing_mode*, *filter_mode* or *normalized_coords* or

combination of these argument values are not valid.

- ✚ CL_INVALID_OPERATION if images are not supported by any device associated with *context* (i.e. CL_DEVICE_IMAGE_SUPPORT specified in *table 4.3* is CL_FALSE).
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

cl_int **clRetainSampler** (cl_sampler *sampler*)

increments the *sampler* reference count. **clCreateSampler** performs an implicit retain. **clRetainSampler** returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_SAMPLER if *sampler* is not a valid sampler object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

cl_int **clReleaseSampler** (cl_sampler *sampler*)

decrements the *sampler* reference count. The sampler object is deleted after the reference count becomes zero and commands queued for execution on a command-queue(s) that use *sampler* have finished. **clReleaseSampler** returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_SAMPLER if *sampler* is not a valid sampler object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.5.2 Sampler Object Queries

The function

```
cl_int          clGetSamplerInfo (cl_sampler sampler,
                                cl_sampler_info param_name,
                                size_t param_value_size,
                                void *param_value,
                                size_t *param_value_size_ret)
```

returns information about the sampler object.

sampler specifies the sampler being queried.

param_name specifies the information to query. The list of supported *param_name* types and the information returned in *param_value* by **clGetSamplerInfo** is described in *table 5.12*.

param_value is a pointer to memory where the appropriate result being queried is returned. If *param_value* is NULL, it is ignored.

param_value_size is used to specify the size in bytes of memory pointed to by *param_value*. This size must be \geq size of return type as described in *table 5.12*.

param_value_size_ret returns the actual size in bytes of data copied to *param_value*. If *param_value_size_ret* is NULL, it is ignored.

| cl_sampler_info | Return Type | Info. returned in <i>param_value</i> |
|---|--------------------|---|
| CL_SAMPLER_REFERENCE_COUNT ¹⁴ | cl_uint | Return the <i>sampler</i> reference count. |
| CL_SAMPLER_CONTEXT | cl_context | Return the context specified when the sampler is created. |
| CL_SAMPLER_NORMALIZED_COORDS | cl_bool | Return the normalized coords value associated with <i>sampler</i> . |
| CL_SAMPLER_ADDRESSING_MODE | cl_addressing_mode | Return the addressing mode value associated with <i>sampler</i> . |
| CL_SAMPLER_FILTER_MODE | cl_filter_mode | Return the filter mode value associated with <i>sampler</i> . |

Table 5.12 *clGetSamplerInfo* parameter queries.

¹⁴ The reference count returned should be considered immediately stale. It is unsuitable for general use in applications. This feature is provided for identifying memory leaks.

clGetSamplerInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_VALUE if *param_name* is not valid, or if size in bytes specified by *param_value_size* is < size of return type as described in *table 5.12* and *param_value* is not NULL.
- ✚ CL_INVALID_SAMPLER if *sampler* is a not a valid sampler object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.6 Program Objects

An OpenCL program consists of a set of kernels that are identified as functions declared with the `__kernel` qualifier in the program source. OpenCL programs may also contain auxiliary functions and constant data that can be used by `__kernel` functions. The program executable can be generated *online* or *offline* by the OpenCL compiler for the appropriate target device(s).

A program object encapsulates the following information:

- ✚ An associated context.
- ✚ A program source or binary.
- ✚ The latest successfully built program executable, library or compiled binary, the list of devices for which the program executable, library or compiled binary is built, the build options used and a build log.
- ✚ The number of kernel objects currently attached.

5.6.1 Creating Program Objects

The function

```
cl_program  clCreateProgramWithSource (cl_context context,
                                       cl_uint count,
                                       const char **strings,
                                       const size_t *lengths,
                                       cl_int *errcode_ret)
```

creates a program object for a context, and loads the source code specified by the text strings in the *strings* array into the program object. The devices associated with the program object are the devices associated with *context*. The source code specified by *strings* is either an OpenCL C program source, header or implementation-defined source for custom devices that support an online compiler.

context must be a valid OpenCL context.

strings is an array of *count* pointers to optionally null-terminated character strings that make up the source code.

The *lengths* argument is an array with the number of chars in each string (the string length). If an element in *lengths* is zero, its accompanying string is null-terminated. If *lengths* is NULL, all strings in the *strings* argument are considered null-terminated. Any length value passed in that is greater than zero excludes the null terminator in its count.

errcode_ret will return an appropriate error code. If *errcode_ret* is NULL, no error code is returned.

clCreateProgramWithSource returns a valid non-zero program object and *errcode_ret* is set to CL_SUCCESS if the program object is created successfully. Otherwise, it returns a NULL value with one of the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_CONTEXT if *context* is not a valid context.
- ✚ CL_INVALID_VALUE if *count* is zero or if *strings* or any entry in *strings* is NULL.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```
cl_program  clCreateProgramWithBinary (cl_context context,
                                     cl_uint num_devices,
                                     const cl_device_id *device_list,
                                     const size_t *lengths,
                                     const unsigned char **binaries,
                                     cl_int *binary_status,
                                     cl_int *errcode_ret)
```

creates a program object for a context, and loads the binary bits specified by *binary* into the program object.

context must be a valid OpenCL context.

device_list is a pointer to a list of devices that are in *context*. *device_list* must be a non-NULL value. The binaries are loaded for devices specified in this list.

num_devices is the number of devices listed in *device_list*.

The devices associated with the program object will be the list of devices specified by *device_list*. The list of devices specified by *device_list* must be devices associated with *context*.

lengths is an array of the size in bytes of the program binaries to be loaded for devices specified by *device_list*.

binaries is an array of pointers to program binaries to be loaded for devices specified by *device_list*. For each device given by *device_list*[*i*], the pointer to the program binary for that

device is given by *binaries*[i] and the length of this corresponding binary is given by *lengths*[i]. *lengths*[i] cannot be zero and *binaries*[i] cannot be a NULL pointer.

The program binaries specified by *binaries* contain the bits that describe one of the following:

- ✚ a program executable to be run on the device(s) associated with *context*,
- ✚ a compiled program for device(s) associated with *context*, or
- ✚ a library of compiled programs for device(s) associated with *context*.

The program binary can consist of either or both:

- ✚ Device-specific code and/or,
- ✚ Implementation-specific intermediate representation (IR) which will be converted to the device-specific code.

binary_status returns whether the program binary for each device specified in *device_list* was loaded successfully or not. It is an array of *num_devices* entries and returns CL_SUCCESS in *binary_status*[i] if binary was successfully loaded for device specified by *device_list*[i]; otherwise returns CL_INVALID_VALUE if *lengths*[i] is zero or if *binaries*[i] is a NULL value or CL_INVALID_BINARY in *binary_status*[i] if program binary is not a valid binary for the specified device. If *binary_status* is NULL, it is ignored.

errcode_ret will return an appropriate error code. If *errcode_ret* is NULL, no error code is returned.

clCreateProgramWithBinary returns a valid non-zero program object and *errcode_ret* is set to CL_SUCCESS if the program object is created successfully. Otherwise, it returns a NULL value with one of the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_CONTEXT if *context* is not a valid context.
- ✚ CL_INVALID_VALUE if *device_list* is NULL or *num_devices* is zero.
- ✚ CL_INVALID_DEVICE if OpenCL devices listed in *device_list* are not in the list of devices associated with *context*.
- ✚ CL_INVALID_VALUE if *lengths* or *binaries* are NULL or if any entry in *lengths*[i] is zero or *binaries*[i] is NULL.
- ✚ CL_INVALID_BINARY if an invalid program binary was encountered for any device. *binary_status* will return specific status for each device.

- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

OpenCL allows applications to create a program object using the program source or binary and build appropriate program executables. This can be very useful as it allows applications to load program source and then compile and link to generate a program executable online on its first instance for appropriate OpenCL devices in the system. These executables can now be queried and cached by the application. Future instances of the application launching will no longer need to compile and link the program executables. The cached executables can be read and loaded by the application, which can help significantly reduce the application initialization time.

The function

```
cl_program  clCreateProgramWithBuiltInKernels (cl_context context,
                                              cl_uint num_devices,
                                              const cl_device_id *device_list,
                                              const char *kernel_names,
                                              cl_int *errcode_ret)
```

creates a program object for a context, and loads the information related to the built-in kernels into a program object.

context must be a valid OpenCL context.

num_devices is the number of devices listed in *device_list*.

device_list is a pointer to a list of devices that are in *context*. *device_list* must be a non-NULL value. The built-in kernels are loaded for devices specified in this list.

The devices associated with the program object will be the list of devices specified by *device_list*. The list of devices specified by *device_list* must be devices associated with *context*.

kernel_names is a semi-colon separated list of built-in kernel names.

clCreateProgramWithBuiltInKernels returns a valid non-zero program object and *errcode_ret* is set to CL_SUCCESS if the program object is created successfully. Otherwise, it returns a NULL value with one of the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_CONTEXT if *context* is not a valid context.
- ✚ CL_INVALID_VALUE if *device_list* is NULL or *num_devices* is zero.

- ✚ CL_INVALID_VALUE if *kernel_names* is NULL or *kernel_names* contains a kernel name that is not supported by any of the devices in *device_list*.
- ✚ CL_INVALID_DEVICE if devices listed in *device_list* are not in the list of devices associated with *context*.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

cl_int **clRetainProgram** (cl_program *program*)

increments the *program* reference count. **clCreateProgram** does an implicit retain. **clRetainProgram** returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_PROGRAM if *program* is not a valid program object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

cl_int **clReleaseProgram** (cl_program *program*)

decrements the *program* reference count. The program object is deleted after all kernel objects associated with *program* have been deleted and the *program* reference count becomes zero. **clReleaseProgram** returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_PROGRAM if *program* is not a valid program object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.6.2 Building Program Executables

The function

```
cl_int      clBuildProgram (cl_program program,
                             cl_uint num_devices,
                             const cl_device_id *device_list,
                             const char *options,
                             void (CL_CALLBACK *pfm_notify)(cl_program program,
                                                             void *user_data),
                             void *user_data)
```

builds (compiles & links) a program executable from the program source or binary for all the devices or a specific device(s) in the OpenCL context associated with *program*. OpenCL allows program executables to be built using the source or the binary. **clBuildProgram** must be called for *program* created using either **clCreateProgramWithSource** or **clCreateProgramWithBinary** to build the program executable for one or more devices associated with *program*. If *program* is created with **clCreateProgramWithBinary**, then the program binary must be an executable binary (not a compiled binary or library).

The executable binary can be queried using **clGetProgramInfo**(*program*, CL_PROGRAM_BINARIES, ...) and can be specified to **clCreateProgramWithBinary** to create a new program object.

program is the program object.

device_list is a pointer to a list of devices associated with *program*. If *device_list* is a NULL value, the program executable is built for all devices associated with *program* for which a source or binary has been loaded. If *device_list* is a non-NULL value, the program executable is built for devices specified in this list for which a source or binary has been loaded.

num_devices is the number of devices listed in *device_list*.

options is a pointer to a null-terminated string of characters that describes the build options to be used for building the program executable. The list of supported options is described in section 5.6.4.

pfm_notify is a function pointer to a notification routine. The notification routine is a callback function that an application can register and which will be called when the program executable has been built (successfully or unsuccessfully). If *pfm_notify* is not NULL, **clBuildProgram** does not need to wait for the build to complete and can return immediately once the build operation can begin. The build operation can begin if the context, program whose sources are being compiled and linked, list of devices and build options specified are all valid and appropriate host and device resources needed to perform the build are available. If *pfm_notify* is NULL, **clBuildProgram** does not return until the build has completed. This callback function may be

called asynchronously by the OpenCL implementation. It is the application's responsibility to ensure that the callback function is thread-safe.

user_data will be passed as an argument when *pfn_notify* is called. *user_data* can be NULL.

clBuildProgram returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_PROGRAM if *program* is not a valid program object.
- ✚ CL_INVALID_VALUE if *device_list* is NULL and *num_devices* is greater than zero, or if *device_list* is not NULL and *num_devices* is zero.
- ✚ CL_INVALID_VALUE if *pfn_notify* is NULL but *user_data* is not NULL.
- ✚ CL_INVALID_DEVICE if OpenCL devices listed in *device_list* are not in the list of devices associated with *program*
- ✚ CL_INVALID_BINARY if *program* is created with **clCreateProgramWithBinary** and devices listed in *device_list* do not have a valid program binary loaded.
- ✚ CL_INVALID_BUILD_OPTIONS if the build options specified by *options* are invalid.
- ✚ CL_INVALID_OPERATION if the build of a program executable for any of the devices listed in *device_list* by a previous call to **clBuildProgram** for *program* has not completed.
- ✚ CL_COMPILER_NOT_AVAILABLE if *program* is created with **clCreateProgramWithSource** and a compiler is not available i.e. CL_DEVICE_COMPILER_AVAILABLE specified in *table 4.3* is set to CL_FALSE.
- ✚ CL_BUILD_PROGRAM_FAILURE if there is a failure to build the program executable. This error will be returned if **clBuildProgram** does not return until the build has completed.
- ✚ CL_INVALID_OPERATION if there are kernel objects attached to *program*.
- ✚ CL_INVALID_OPERATION if *program* was not created with **clCreateProgramWithSource** or **clCreateProgramWithBinary**.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.6.3 Separate Compilation and Linking of Programs

OpenCL 1.2 extends how programs are compiled and linked to support the following:

- ✚ Separate compilation and link stages. Program sources can be compiled to generate a compiled binary object and linked in a separate stage with other compiled program objects to the program executable.
- ✚ Embedded headers. In OpenCL 1.0 and 1.1, the `-I` build option could be used to specify the list of directories to be searched for headers files that are included by a program source(s). OpenCL 1.2 extends this by allowing the header sources to come from program objects instead of just header files.
- ✚ Libraries. The linker can be used to link compiled objects and libraries into a program executable or to create a library of compiled binaries.

The function

```
cl_int      clCompileProgram (cl_program program,
                               cl_uint num_devices,
                               const cl_device_id *device_list,
                               const char *options,
                               cl_uint num_input_headers,
                               const cl_program *input_headers,
                               const char **header_include_names,
                               void (CL_CALLBACK *pfm_notify)(cl_program program,
                                                             void *user_data),
                               void *user_data)
```

compiles a program's source for all the devices or a specific device(s) in the OpenCL context associated with *program*. The pre-processor runs before the program sources are compiled. The compiled binary is built for all devices associated with *program* or the list of devices specified. The compiled binary can be queried using **clGetProgramInfo**(*program*, CL_PROGRAM_BINARIES, ...) and can be specified to **clCreateProgramWithBinary** to create a new program object.

program is the program object that is the compilation target.

device_list is a pointer to a list of devices associated with *program*. If *device_list* is a NULL value, the compile is performed for all devices associated with *program*. If *device_list* is a non-NULL value, the compile is performed for devices specified in this list.

num_devices is the number of devices listed in *device_list*.

options is a pointer to a null-terminated string of characters that describes the compilation options to be used for building the program executable. The list of supported options is as described in *section 5.6.4*.

num_input_headers specifies the number of programs that describe headers in the array referenced by *input_headers*.

input_headers is an array of program embedded headers created with **clCreateProgramWithSource**.

header_include_names is an array that has a one to one correspondence with *input_headers*. Each entry in *header_include_names* specifies the include name used by source in *program* that comes from an embedded header. The corresponding entry in *input_headers* identifies the program object which contains the header source to be used. The embedded headers are first searched before the headers in the list of directories specified by the `-I` compile option (as described in *section 5.6.4.1*). If multiple entries in *header_include_names* refer to the same header name, the first one encountered will be used.

For example, consider the following program source:

```
#include <foo.h>
#include <mydir/myinc.h>

__kernel void
image_filter (int n, int m,
              __constant float *filter_weights,
              __read_only image2d_t src_image,
              __write_only image2d_t dst_image)
{
    ...
}
```

This kernel includes two headers `foo.h` and `mydir/myinc.h`. The following describes how these headers can be passed as embedded headers in program objects:

```
cl_program foo_pg = clCreateProgramWithSource(context,
                                             1, &foo_header_src, NULL, &err);
cl_program myinc_pg = clCreateProgramWithSource(context,
                                                1, &myinc_header_src, NULL, &err);

// let's assume the program source described above is given
// by program_A and is loaded via clCreateProgramWithSource

cl_program input_headers[2] = { foo_pg, myinc_pg };
char * input_header_names[2] = { "foo.h", "mydir/myinc.h" };
clCompileProgram(program_A,
```

```

0, NULL, // num_devices & device_list
NULL, // compile_options
2, // num_input_headers
input_headers,
input_header_names,
NULL, NULL); // pfn_notify & user_data

```

pfn_notify is a function pointer to a notification routine. The notification routine is a callback function that an application can register and which will be called when the program executable has been built (successfully or unsuccessfully). If *pfn_notify* is not NULL, **clCompileProgram** does not need to wait for the compiler to complete and can return immediately once the compilation can begin. The compilation can begin if the context, program whose sources are being compiled, list of devices, input headers, programs that describe input headers and compiler options specified are all valid and appropriate host and device resources needed to perform the compile are available. If *pfn_notify* is NULL, **clCompileProgram** does not return until the compiler has completed. This callback function may be called asynchronously by the OpenCL implementation. It is the application's responsibility to ensure that the callback function is thread-safe.

user_data will be passed as an argument when *pfn_notify* is called. *user_data* can be NULL.

clCompileProgram returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_PROGRAM if *program* is not a valid program object.
- ✚ CL_INVALID_VALUE if *device_list* is NULL and *num_devices* is greater than zero, or if *device_list* is not NULL and *num_devices* is zero.
- ✚ CL_INVALID_VALUE if *num_input_headers* is zero and *header_include_names* or *input_headers* are not NULL or if *num_input_headers* is not zero and *header_include_names* or *input_headers* are NULL.
- ✚ CL_INVALID_VALUE if *pfn_notify* is NULL but *user_data* is not NULL.
- ✚ CL_INVALID_DEVICE if OpenCL devices listed in *device_list* are not in the list of devices associated with *program*
- ✚ CL_INVALID_COMPILER_OPTIONS if the compiler options specified by *options* are invalid.
- ✚ CL_INVALID_OPERATION if the compilation or build of a program executable for any of the devices listed in *device_list* by a previous call to **clCompileProgram** or **clBuildProgram** for *program* has not completed.
- ✚ CL_COMPILER_NOT_AVAILABLE if a compiler is not available i.e.

CL_DEVICE_COMPILER_AVAILABLE specified in *table 4.3* is set to CL_FALSE.

- ✚ CL_COMPILE_PROGRAM_FAILURE if there is a failure to compile the program source. This error will be returned if **clCompileProgram** does not return until the compile has completed.
- ✚ CL_INVALID_OPERATION if there are kernel objects attached to *program*.
- ✚ CL_INVALID_OPERATION if *program* has no source i.e. it has not been created with **clCreateProgramWithSource**.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```
cl_program  clLinkProgram (cl_context context,
                           cl_uint num_devices,
                           const cl_device_id *device_list,
                           const char *options,
                           cl_uint num_input_programs,
                           const cl_program *input_programs,
                           void (CL_CALLBACK *pfm_notify)(cl_program program,
                                                           void *user_data),
                           void *user_data,
                           cl_int *errcode_ret)
```

links a set of compiled program objects and libraries for all the devices or a specific device(s) in the OpenCL context and creates an executable. **clLinkProgram** creates a new program object which contains this executable. The executable binary can be queried using **clGetProgramInfo**(*program*, CL_PROGRAM_BINARIES, ...) and can be specified to **clCreateProgramWithBinary** to create a new program object.

The devices associated with the returned program object will be the list of devices specified by *device_list* or if *device_list* is NULL it will be the list of devices associated with *context*.

context must be a valid OpenCL context.

device_list is a pointer to a list of devices that are in *context*. If *device_list* is a NULL value, the link is performed for all devices associated with *context* for which a compiled object is available. If *device_list* is a non-NULL value, the link is performed for devices specified in this list for which a compiled object is available.

num_devices is the number of devices listed in *device_list*.

options is a pointer to a null-terminated string of characters that describes the link options to be used for building the program executable. The list of supported options is as described in section 5.6.5.

num_input_programs specifies the number of programs in array referenced by *input_programs*.

input_programs is an array of program objects that are compiled binaries or libraries that are to be linked to create the program executable. For each device in *device_list* or if *device_list* is NULL the list of devices associated with context, the following cases occur:

- ✚ All programs specified by *input_programs* contain a compiled binary or library for the device. In this case, a link is performed to generate a program executable for this device.
- ✚ None of the programs contain a compiled binary or library for that device. In this case, no link is performed and there will be no program executable generated for this device.
- ✚ All other cases will return a CL_INVALID_OPERATION error.

pfn_notify is a function pointer to a notification routine. The notification routine is a callback function that an application can register and which will be called when the program executable has been built (successfully or unsuccessfully).

If *pfn_notify* is not NULL, **clLinkProgram** does not need to wait for the linker to complete and can return immediately once the linking operation can begin. Once the linker has completed, the *pfn_notify* callback function is called which returns the program object returned by **clLinkProgram**. The application can query the link status and log for this program object. This callback function may be called asynchronously by the OpenCL implementation. It is the application's responsibility to ensure that the callback function is thread-safe.

If *pfn_notify* is NULL, **clLinkProgram** does not return until the linker has completed.

user_data will be passed as an argument when *pfn_notify* is called. *user_data* can be NULL.

The linking operation can begin if the context, list of devices, input programs and linker options specified are all valid and appropriate host and device resources needed to perform the link are available. If the linking operation can begin, **clLinkProgram** returns a valid non-zero program object.

If *pfn_notify* is NULL, the *errcode_ret* will be set to CL_SUCCESS if the link operation was successful and CL_LINK_FAILURE if there is a failure to link the compiled binaries and/or libraries.

If *pfn_notify* is not NULL, **clLinkProgram** does not have to wait until the linker to complete and

can return `CL_SUCCESS` in *errcode_ret* if the linking operation can begin. The *pfm_notify* callback function will return a `CL_SUCCESS` or `CL_LINK_FAILURE` if the linking operation was successful or not.

Otherwise **clLinkProgram** returns a NULL program object with an appropriate error in *errcode_ret*. The application should query the linker status of this program object to check if the link was successful or not. The list of errors that can be returned are:

- ✚ `CL_INVALID_CONTEXT` if *context* is not a valid context.
- ✚ `CL_INVALID_VALUE` if *device_list* is NULL and *num_devices* is greater than zero, or if *device_list* is not NULL and *num_devices* is zero.
- ✚ `CL_INVALID_VALUE` if *num_input_programs* is zero and *input_programs* is NULL or if *num_input_programs* is zero and *input_programs* is not NULL or if *num_input_programs* is not zero and *input_programs* is NULL.
- ✚ `CL_INVALID_PROGRAM` if programs specified in *input_programs* are not valid program objects.
- ✚ `CL_INVALID_VALUE` if *pfm_notify* is NULL but *user_data* is not NULL.
- ✚ `CL_INVALID_DEVICE` if OpenCL devices listed in *device_list* are not in the list of devices associated with *context*
- ✚ `CL_INVALID_LINKER_OPTIONS` if the linker options specified by *options* are invalid.
- ✚ `CL_INVALID_OPERATION` if the compilation or build of a program executable for any of the devices listed in *device_list* by a previous call to **clCompileProgram** or **clBuildProgram** for *program* has not completed.
- ✚ `CL_INVALID_OPERATION` if the rules for devices containing compiled binaries or libraries as described in *input_programs* argument above are not followed.
- ✚ `CL_LINKER_NOT_AVAILABLE` if a linker is not available i.e. `CL_DEVICE_LINKER_AVAILABLE` specified in *table 4.3* is set to `CL_FALSE`.
- ✚ `CL_LINK_PROGRAM_FAILURE` if there is a failure to link the compiled binaries and/or libraries.
- ✚ `CL_OUT_OF_RESOURCES` if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ `CL_OUT_OF_HOST_MEMORY` if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.6.4 Compiler Options

The compiler options are categorized as pre-processor options, options for math intrinsics, options that control optimization and miscellaneous options. This specification defines a standard set of options that must be supported by the OpenCL C compiler when building program executables online or offline. These may be extended by a set of vendor- or platform-specific options.

5.6.4.1 Preprocessor options

These options control the OpenCL C preprocessor which is run on each program source before actual compilation.

-D *name*

Predefine *name* as a macro, with definition 1.

-D *name=definition*

The contents of *definition* are tokenized and processed as if they appeared during translation phase three in a '#define' directive. In particular, the definition will be truncated by embedded newline characters.

-D options are processed in the order they are given in the *options* argument to **clBuildProgram** or **clCompileProgram**.

-I *dir*

Add the directory *dir* to the list of directories to be searched for header files.

5.6.4.2 Math Intrinsic Options

These options control compiler behavior regarding floating-point arithmetic. These options trade off between speed and correctness.

-cl-single-precision-constant

Treat double precision floating-point constant as single precision constant.

-cl-denorms-are-zero

This option controls how single precision and double precision denormalized numbers are handled. If specified as a build option, the single precision denormalized numbers may be flushed to zero; double precision denormalized numbers may also be flushed to zero if the optional extension for double precision is supported. This is intended to be a performance hint and the OpenCL compiler can choose not to flush denorms to zero if the device supports single precision (or double precision) denormalized numbers.

This option is ignored for single precision numbers if the device does not support single precision denormalized numbers i.e. `CL_FP_DENORM` bit is not set in `CL_DEVICE_SINGLE_FP_CONFIG`.

This option is ignored for double precision numbers if the device does not support double precision or if it does support double precision but not double precision denormalized numbers i.e. `CL_FP_DENORM` bit is not set in `CL_DEVICE_DOUBLE_FP_CONFIG`.

This flag only applies for scalar and vector single precision floating-point variables and computations on these floating-point variables inside a program. It does not apply to reading from or writing to image objects.

`-cl-fp32-correctly-rounded-divide-sqrt`

The `-cl-fp32-correctly-rounded-divide-sqrt` build option to **clBuildProgram** or **clCompileProgram** allows an application to specify that single precision floating-point divide (`x/y` and `1/x`) and `sqrt` used in the program source are correctly rounded. If this build option is not specified, the minimum numerical accuracy of single precision floating-point divide and `sqrt` are as defined in *section 7.4* of the OpenCL specification.

This build option can only be specified if the `CL_FP_CORRECTLY_ROUNDED_DIVIDE_SQRT` is set in `CL_DEVICE_SINGLE_FP_CONFIG` (as defined in *table 4.3*) for devices that the program is being build. **clBuildProgram** or **clCompileProgram** will fail to compile the program for a device if the `-cl-fp32-correctly-rounded-divide-sqrt` option is specified and `CL_FP_CORRECTLY_ROUNDED_DIVIDE_SQRT` is not set for the device.

5.6.4.3 Optimization Options

These options control various sorts of optimizations. Turning on optimization flags makes the compiler attempt to improve the performance and/or code size at the expense of compilation time and possibly the ability to debug the program.

`-cl-opt-disable`

This option disables all optimizations. The default is optimizations are enabled.

The following options control compiler behavior regarding floating-point arithmetic. These options trade off between performance and correctness and must be specifically enabled. These options are not turned on by default since it can result in incorrect output for programs which depend on an exact implementation of IEEE 754 rules/specifications for math functions.

`-cl-mad-enable`

Allow `a * b + c` to be replaced by `a mad`. The `mad` computes `a * b + c` with reduced accuracy. For example, some OpenCL devices implement `mad` as truncate the result of `a * b` before adding it to `c`.

-cl-no-signed-zeros

Allow optimizations for floating-point arithmetic that ignore the signedness of zero. IEEE 754 arithmetic specifies the distinct behavior of $+0.0$ and -0.0 values, which then prohibits simplification of expressions such as $x+0.0$ or $0.0*x$ (even with `-cl-finite-math` only). This option implies that the sign of a zero result isn't significant.

-cl-unsafe-math-optimizations

Allow optimizations for floating-point arithmetic that (a) assume that arguments and results are valid, (b) may violate IEEE 754 standard and (c) may violate the OpenCL numerical compliance requirements as defined in *section 7.4* for single precision and double precision floating-point, and edge case behavior in *section 7.5*. This option includes the `-cl-no-signed-zeros` and `-cl-mad-enable` options.

-cl-finite-math-only

Allow optimizations for floating-point arithmetic that assume that arguments and results are not NaNs or $\pm\infty$. This option may violate the OpenCL numerical compliance requirements defined in *section 7.4* for single precision and double precision floating-point, and edge case behavior in *section 7.5*.

-cl-fast-relaxed-math

Sets the optimization options `-cl-finite-math-only` and `-cl-unsafe-math-optimizations`. This allows optimizations for floating-point arithmetic that may violate the IEEE 754 standard and the OpenCL numerical compliance requirements defined in *section 7.4* for single precision and double precision floating-point, and edge case behavior in *section 7.5*. This option causes the preprocessor macro `__FAST_RELAXED_MATH__` to be defined in the OpenCL program.

5.6.4.4 Options to Request or Suppress Warnings

Warnings are diagnostic messages that report constructions which are not inherently erroneous but which are risky or suggest there may have been an error. The following language-independent options do not enable specific warnings but control the kinds of diagnostics produced by the OpenCL compiler.

-w

Inhibit all warning messages.

-Werror

Make all warnings into errors.

5.6.4.5 Options Controlling the OpenCL C version

The following option controls the version of OpenCL C that the compiler accepts.

-cl-std=

Determine the OpenCL C language version to use. A value for this option must be provided. Valid values are:

CL1.1 – Support all OpenCL C programs that use the OpenCL C language features defined in *section 6* of the OpenCL 1.1 specification.

CL1.2 – Support all OpenCL C programs that use the OpenCL C language features defined in *section 6* of the OpenCL 1.2 specification.

Calls to **clBuildProgram** or **clCompileProgram** with the -cl-std=CL1.1 option **will fail** to compile the program for any devices with CL_DEVICE_OPENCL_C_VERSION = OpenCL C 1.0. Calls to **clBuildProgram** or **clCompileProgram** with the -cl-std=CL1.2 option **will fail** to compile the program for any devices with CL_DEVICE_OPENCL_C_VERSION = OpenCL C 1.0 or OpenCL C 1.1.

If the -cl-std build option is not specified, the CL_DEVICE_OPENCL_C_VERSION is used to select the version of OpenCL C to be used when compiling the program for each device.

5.6.4.6 Options for Querying Kernel Argument Information

-cl-kernel-arg-info

This option allows the compiler to store information about the arguments of a kernel(s) in the program executable. The argument information stored includes the argument name, its type, the address and access qualifiers used. Refer to description of **clGetKernelArgInfo** on how to query this information.

5.6.5 Linker Options

This specification defines a standard set of linker options that must be supported by the OpenCL C compiler when linking compiled programs online or offline. These linker options are categorized as library linking options and program linking options. These may be extended by a set of vendor- or platform-specific options.

5.6.5.1 Library Linking Options

The following options can be specified when creating a library of compiled binaries.

-create-library

Create a library of compiled binaries specified in *input_programs* argument to **clLinkProgram**.

-enable-link-options

Allows the linker to modify the library behavior based on one or more link options (described in *section 5.6.5.2*) when this library is linked with a program executable. This option must be specified with the `–create-library` option.

5.6.5.2 Program Linking Options

The following options can be specified when linking a program executable.

-cl-denorms-are-zero

-cl-no-signed-zeroes

-cl-unsafe-math-optimizations

-cl-finite-math-only

-cl-fast-relaxed-math

The options are described in *section 5.6.4.2* and *section 5.6.4.3*. The linker may apply these options to all compiled program objects specified to **clLinkProgram**. The linker may apply these options only to libraries which were created with the `–enable-link-option`.


5.6.6 Unloading the OpenCL Compiler

The function

```
cl_int  clUnloadPlatformCompiler (cl_platform_id platform)
```

allows the implementation to release the resources allocated by the OpenCL compiler for *platform*. This is a hint from the application and does not guarantee that the compiler will not be used in the future or that the compiler will actually be unloaded by the implementation. Calls to **clBuildProgram**, **clCompileProgram** or **clLinkProgram** after **clUnloadPlatformCompiler** will reload the compiler, if necessary, to build the appropriate program executable.

clUnloadPlatformCompiler returns `CL_SUCCESS` if the function is executed successfully. Otherwise, it returns one of the following errors:

 `CL_INVALID_PLATFORM` if *platform* is not a valid platform.

5.6.7 Program Object Queries

The function

```
cl_int          clGetProgramInfo (cl_program program,
                                cl_program_info param_name,
                                size_t param_value_size,
                                void *param_value,
                                size_t *param_value_size_ret)
```

returns information about the program object.

program specifies the program object being queried.

param_name specifies the information to query. The list of supported *param_name* types and the information returned in *param_value* by **clGetProgramInfo** is described in *table 5.13*.

param_value is a pointer to memory where the appropriate result being queried is returned. If *param_value* is NULL, it is ignored.

param_value_size is used to specify the size in bytes of memory pointed to by *param_value*. This size must be \geq size of return type as described in *table 5.13*.

param_value_size_ret returns the actual size in bytes of data copied to *param_value*. If *param_value_size_ret* is NULL, it is ignored.

| cl_program_info | Return Type | Info. returned in <i>param_value</i> |
|--|----------------|--|
| CL_PROGRAM_REFERENCE_COUNT ¹⁵ | cl_uint | Return the <i>program</i> reference count. |
| CL_PROGRAM_CONTEXT | cl_context | Return the context specified when the program object is created |
| CL_PROGRAM_NUM_DEVICES | cl_uint | Return the number of devices associated with <i>program</i> . |
| CL_PROGRAM_DEVICES | cl_device_id[] | Return the list of devices associated with the program object. This can be the devices associated with context on which the program object has been created or can be a subset of devices that are specified when a program object is created using clCreateProgramWithBinary . |

¹⁵ The reference count returned should be considered immediately stale. It is unsuitable for general use in applications. This feature is provided for identifying memory leaks.

| | | |
|--------------------------------|-------------------|--|
| CL_PROGRAM_SOURCE | char[] | <p>Return the program source code specified by clCreateProgramWithSource. The source string returned is a concatenation of all source strings specified to clCreateProgramWithSource with a null terminator. The concatenation strips any nulls in the original source strings.</p> <p>If <i>program</i> is created using clCreateProgramWithBinary or clCreateProgramWithBuiltinKernels, a null string or the appropriate program source code is returned depending on whether or not the program source code is stored in the binary.</p> <p>The actual number of characters that represents the program source code including the null terminator is returned in <i>param_value_size_ret</i>.</p> |
| CL_PROGRAM_BINARY_SIZES | size_t[] | <p>Returns an array that contains the size in bytes of the program binary (could be an executable binary, compiled binary or library binary) for each device associated with <i>program</i>. The size of the array is the number of devices associated with <i>program</i>. If a binary is not available for a device(s), a size of zero is returned.</p> <p>If <i>program</i> is created using clCreateProgramWithBuiltinKernels, the implementation may return zero in any entries of the returned array.</p> |
| CL_PROGRAM_BINARIES | unsigned char *[] | <p>Return the program binaries (could be an executable binary, compiled binary or library binary) for all devices associated with <i>program</i>. For each device in <i>program</i>, the binary returned can be the binary specified for the device when <i>program</i> is created with clCreateProgramWithBinary or it can be the executable binary generated by clBuildProgram or clLinkProgram. If <i>program</i> is created with</p> |

| | | |
|-------------------------------|--------|---|
| | | <p>clCreateProgramWithSource, the binary returned is the binary generated by clBuildProgram, clCompileProgram or clLinkProgram. The bits returned can be an implementation-specific intermediate representation (a.k.a. IR) or device specific executable bits or both. The decision on which information is returned in the binary is up to the OpenCL implementation.</p> <p><i>param_value</i> points to an array of <i>n</i> pointers allocated by the caller, where <i>n</i> is the number of devices associated with program. The buffer sizes needed to allocate the memory that these <i>n</i> pointers refer to can be queried using the CL_PROGRAM_BINARY_SIZES query as described in this table.</p> <p>Each entry in this array is used by the implementation as the location in memory where to copy the program binary for a specific device, if there is a binary available. To find out which device the program binary in the array refers to, use the CL_PROGRAM_DEVICES query to get the list of devices. There is a one-to-one correspondence between the array of <i>n</i> pointers returned by CL_PROGRAM_BINARIES and array of devices returned by CL_PROGRAM_DEVICES.</p> <p>If an entry value in the array is NULL, the implementation skips copying the program binary for the specific device identified by the array index.</p> |
| | | |
| CL_PROGRAM_NUM_KERNELS | size_t | Returns the number of kernels declared in <i>program</i> that can be created with clCreateKernel . This information is only available after a successful program executable has been built for at least one device in the list of devices associated with <i>program</i> . |
| CL_PROGRAM_KERNEL_ | char[] | Returns a semi-colon separated list of |

| | | |
|-------|--|---|
| NAMES | | kernel names in <i>program</i> that can be created with clCreateKernel . This information is only available after a successful program executable has been built for at least one device in the list of devices associated with <i>program</i> . |
|-------|--|---|

Table 5.13 *clGetProgramInfo* parameter queries.

clGetProgramInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_VALUE if *param_name* is not valid, or if size in bytes specified by *param_value_size* is < size of return type as described in *table 5.13* and *param_value* is not NULL.
- ✚ CL_INVALID_PROGRAM if *program* is a not a valid program object.
- ✚ CL_INVALID_PROGRAM_EXECUTABLE if *param_name* is CL_PROGRAM_NUM_KERNELS or CL_PROGRAM_KERNEL_NAMES and a successful program executable has not been built for at least one device in the list of devices associated with *program*.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```

cl_int          clGetProgramBuildInfo (cl_program program,
                                         cl_device_id device,
                                         cl_program_build_info param_name,
                                         size_t param_value_size,
                                         void *param_value,
                                         size_t *param_value_size_ret)

```

returns build information for each device in the program object.

program specifies the program object being queried.

device specifies the device for which build information is being queried. *device* must be a valid device associated with *program*.

param_name specifies the information to query. The list of supported *param_name* types and the information returned in *param_value* by **clGetProgramBuildInfo** is described in *table 5.14*.

param_value is a pointer to memory where the appropriate result being queried is returned. If *param_value* is NULL, it is ignored.

param_value_size is used to specify the size in bytes of memory pointed to by *param_value*. This size must be \geq size of return type as described in *table 5.14*.

param_value_size_ret returns the actual size in bytes of data copied to *param_value*. If *param_value_size_ret* is NULL, it is ignored.

| cl_program_build_info | Return Type | Info. returned in <i>param_value</i> |
|---------------------------------|--------------------|---|
| CL_PROGRAM_BUILD_STATUS | cl_build_status | <p>Returns the build, compile or link status, whichever was performed last on <i>program</i> for <i>device</i>.</p> <p>This can be one of the following:</p> <p>CL_BUILD_NONE. The build status returned if no clBuildProgram, clCompileProgram or clLinkProgram has been performed on the specified program object for <i>device</i>.</p> <p>CL_BUILD_ERROR. The build status returned if clBuildProgram, clCompileProgram or clLinkProgram whichever was performed last on the specified program object for <i>device</i> generated an error.</p> <p>CL_BUILD_SUCCESS. The build status returned if clBuildProgram, clCompileProgram or clLinkProgram whichever was performed last on the specified program object for <i>device</i> was successful.</p> <p>CL_BUILD_IN_PROGRESS. The build status returned if clBuildProgram, clCompileProgram or clLinkProgram whichever was performed last on the specified program object for <i>device</i> has not finished.</p> |
| CL_PROGRAM_BUILD_OPTIONS | char[] | <p>Return the build, compile or link options specified by the <i>options</i> argument in clBuildProgram, clCompileProgram or clLinkProgram, whichever was performed last</p> |

| | | |
|-------------------------------|------------------------|---|
| | | <p>on <i>program</i> for <i>device</i>.</p> <p>If build status of <i>program</i> for <i>device</i> is CL_BUILD_NONE, an empty string is returned.</p> |
| CL_PROGRAM_BUILD_LOG | char[] | <p>Return the build, compile or link log for clBuildProgram, clCompileProgram or clLinkProgram whichever was performed last on <i>program</i> for <i>device</i>.</p> <p>If build status of <i>program</i> for <i>device</i> is CL_BUILD_NONE, an empty string is returned.</p> |
| CL_PROGRAM_BINARY_TYPE | cl_program_binary_type | <p>Return the program binary type for <i>device</i>. This can be one of the following values:</p> <p>CL_PROGRAM_BINARY_TYPE_NONE – There is no binary associated with <i>device</i>.</p> <p>CL_PROGRAM_BINARY_TYPE_COMPILED_OBJECT – A compiled binary is associated with <i>device</i>. This is the case if <i>program</i> was created using clCreateProgramWithSource and compiled using clCompileProgram or a compiled binary is loaded using clCreateProgramWithBinary.</p> <p>CL_PROGRAM_BINARY_TYPE_LIBRARY – A library binary is associated with <i>device</i>. This is the case if <i>program</i> was created by clLinkProgram which is called with the –create-library link option or if a library binary is loaded using clCreateProgramWithBinary.</p> <p>CL_PROGRAM_BINARY_TYPE_EXECUTABLE – An executable binary is associated with <i>device</i>. This is the case if <i>program</i> was created by clLinkProgram without the –create-library link option or program was created by clBuildProgram or an executable binary is loaded using clCreateProgramWithBinary.</p> |

Table 5.14 *clGetProgramBuildInfo* parameter queries.

clGetProgramBuildInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_DEVICE if *device* is not in the list of devices associated with *program*.
- ✚ CL_INVALID_VALUE if *param_name* is not valid, or if size in bytes specified by *param_value_size* is < size of return type as described in *table 5.14* and *param_value* is not NULL.
- ✚ CL_INVALID_PROGRAM if *program* is a not a valid program object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

NOTE:

A program binary (compiled binary, library binary or executable binary) built for a parent device can be used by all its sub-devices. If a program binary has not been built for a sub-device, the program binary associated with the parent device will be used.

A program binary for a device specified with **clCreateProgramWithBinary** or queried using **clGetProgramInfo** can be used as the binary for the associated root device, and all sub-devices created from the root-level device or sub-devices thereof.

5.7 Kernel Objects

A kernel is a function declared in a program. A kernel is identified by the `__kernel` qualifier applied to any function in a program. A kernel object encapsulates the specific `__kernel` function declared in a program and the argument values to be used when executing this `__kernel` function.

5.7.1 Creating Kernel Objects

To create a kernel object, use the function

```
cl_kernel      clCreateKernel (cl_program program,
                               const char *kernel_name,
                               cl_int *errcode_ret)
```

program is a program object with a successfully built executable.

kernel_name is a function name in the program declared with the `__kernel` qualifier.

errcode_ret will return an appropriate error code. If *errcode_ret* is NULL, no error code is returned.

clCreateKernel returns a valid non-zero kernel object and *errcode_ret* is set to CL_SUCCESS if the kernel object is created successfully. Otherwise, it returns a NULL value with one of the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_PROGRAM if *program* is not a valid program object.
- ✚ CL_INVALID_PROGRAM_EXECUTABLE if there is no successfully built executable for *program*.
- ✚ CL_INVALID_KERNEL_NAME if *kernel_name* is not found in *program*.
- ✚ CL_INVALID_KERNEL_DEFINITION if the function definition for `__kernel` function given by *kernel_name* such as the number of arguments, the argument types are not the same for all devices for which the *program* executable has been built.
- ✚ CL_INVALID_VALUE if *kernel_name* is NULL.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the

OpenCL implementation on the host.

The function

```
cl_int      clCreateKernelsInProgram (cl_program program,
                                       cl_uint num_kernels,
                                       cl_kernel *kernels,
                                       cl_uint *num_kernels_ret)
```

creates kernel objects for all kernel functions in *program*. Kernel objects are not created for any `__kernel` functions in *program* that do not have the same function definition across all devices for which a program executable has been successfully built.

program is a program object with a successfully built executable.

num_kernels is the size of memory pointed to by *kernels* specified as the number of `cl_kernel` entries.

kernels is the buffer where the kernel objects for kernels in *program* will be returned. If *kernels* is NULL, it is ignored. If *kernels* is not NULL, *num_kernels* must be greater than or equal to the number of kernels in *program*.

num_kernels_ret is the number of kernels in *program*. If *num_kernels_ret* is NULL, it is ignored.

clCreateKernelsInProgram will return `CL_SUCCESS` if the kernel objects were successfully allocated. Otherwise, it returns one of the following errors:

- ✚ `CL_INVALID_PROGRAM` if *program* is not a valid program object.
- ✚ `CL_INVALID_PROGRAM_EXECUTABLE` if there is no successfully built executable for any device in *program*.
- ✚ `CL_INVALID_VALUE` if *kernels* is not NULL and *num_kernels* is less than the number of kernels in *program*.
- ✚ `CL_OUT_OF_RESOURCES` if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ `CL_OUT_OF_HOST_MEMORY` if there is a failure to allocate resources required by the OpenCL implementation on the host.

Kernel objects can only be created once you have a program object with a valid program source or binary loaded into the program object and the program executable has been successfully built for one or more devices associated with program. No changes to the program executable are allowed while there are kernel objects associated with a program object. This means that calls to

clBuildProgram and **clCompileProgram** return `CL_INVALID_OPERATION` if there are kernel objects attached to a program object. The OpenCL context associated with *program* will be the context associated with *kernel*. The list of devices associated with *program* are the devices associated with *kernel*. Devices associated with a program object for which a valid program executable has been built can be used to execute kernels declared in the program object.

The function

`cl_int` **clRetainKernel** (`cl_kernel kernel`)

increments the *kernel* reference count. **clRetainKernel** returns `CL_SUCCESS` if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ `CL_INVALID_KERNEL` if *kernel* is not a valid kernel object.
- ✚ `CL_OUT_OF_RESOURCES` if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ `CL_OUT_OF_HOST_MEMORY` if there is a failure to allocate resources required by the OpenCL implementation on the host.

clCreateKernel or **clCreateKernelsInProgram** do an implicit retain.

The function

`cl_int` **clReleaseKernel** (`cl_kernel kernel`)

decrements the *kernel* reference count. **clReleaseKernel** returns `CL_SUCCESS` if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ `CL_INVALID_KERNEL` if *kernel* is not a valid kernel object.
- ✚ `CL_OUT_OF_RESOURCES` if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ `CL_OUT_OF_HOST_MEMORY` if there is a failure to allocate resources required by the OpenCL implementation on the host.

The kernel object is deleted once the number of instances that are retained to *kernel* become zero and the kernel object is no longer needed by any enqueued commands that use *kernel*.

5.7.2 Setting Kernel Arguments

To execute a kernel, the kernel arguments must be set.

The function

```
cl_int clSetKernelArg (cl_kernel kernel,
                      cl_uint arg_index,
                      size_t arg_size,
                      const void *arg_value)
```

is used to set the argument value for a specific argument of a kernel.

kernel is a valid kernel object.

arg_index is the argument index. Arguments to the kernel are referred by indices that go from 0 for the leftmost argument to $n - 1$, where n is the total number of arguments declared by a kernel.

For example, consider the following kernel:

```
kernel void
image_filter (int n, int m,
              __constant float *filter_weights,
              __read_only image2d_t src_image,
              __write_only image2d_t dst_image)
{
    ...
}
```

Argument index values for `image_filter` will be 0 for `n`, 1 for `m`, 2 for `filter_weights`, 3 for `src_image` and 4 for `dst_image`.

arg_value is a pointer to data that should be used as the argument value for argument specified by *arg_index*. The argument data pointed to by *arg_value* is copied and the *arg_value* pointer can therefore be reused by the application after `clSetKernelArg` returns. The argument value specified is the value used by all API calls that enqueue *kernel* (`clEnqueueNDRangeKernel` and `clEnqueueTask`) until the argument value is changed by a call to `clSetKernelArg` for *kernel*.

If the argument is a memory object (buffer, image or image array), the *arg_value* entry will be a pointer to the appropriate buffer, image or image array object. The memory object must be created with the context associated with the kernel object. If the argument is a buffer object, the *arg_value* pointer can be NULL or point to a NULL value in which case a NULL value will be used as the value for the argument declared as a pointer to `__global` or `__constant` memory in the kernel. If the argument is declared with the `__local` qualifier, the *arg_value* entry must be NULL. If the argument is of type `sampler_t`, the *arg_value* entry must be a pointer to the sampler object.

If the argument is declared to be a pointer of a built-in scalar or vector type, or a user defined structure type in the global or constant address space, the memory object specified as argument value must be a buffer object (or NULL). If the argument is declared with the `__constant` qualifier, the size in bytes of the memory object cannot exceed `CL_DEVICE_MAX_CONSTANT_BUFFER_SIZE` and the number of arguments declared as pointers to `__constant` memory cannot exceed `CL_DEVICE_MAX_CONSTANT_ARGS`.

The memory object specified as argument value must be a 2D image object if the argument is declared to be of type `image2d_t`. The memory object specified as argument value must be a 3D image object if argument is declared to be of type `image3d_t`. The memory object specified as argument value must be a 1D image object if the argument is declared to be of type `image1d_t`. The memory object specified as argument value must be a 1D image buffer object if the argument is declared to be of type `image1d_buffer_t`. The memory object specified as argument value must be a 1D image array object if argument is declared to be of type `image1d_array_t`. The memory object specified as argument value must be a 2D image array object if argument is declared to be of type `image2d_array_t`.

For all other kernel arguments, the `arg_value` entry must be a pointer to the actual data to be used as argument value.

`arg_size` specifies the size of the argument value. If the argument is a memory object, the size is the size of the buffer or image object type. For arguments declared with the `__local` qualifier, the size specified will be the size in bytes of the buffer that must be allocated for the `__local` argument. If the argument is of type `sampler_t`, the `arg_size` value must be equal to `sizeof(cl_sampler)`. For all other arguments, the size will be the size of argument type.

NOTE: A kernel object does not update the reference count for objects such as memory, sampler objects specified as argument values by `clSetKernelArg`. Users may not rely on a kernel object to retain objects specified as argument values to the kernel¹⁶.

`clSetKernelArg` returns `CL_SUCCESS` if the function was executed successfully. Otherwise, it returns one of the following errors:

- ✚ `CL_INVALID_KERNEL` if `kernel` is not a valid kernel object.
- ✚ `CL_INVALID_ARG_INDEX` if `arg_index` is not a valid argument index.
- ✚ `CL_INVALID_ARG_VALUE` if `arg_value` specified is not a valid value.

¹⁶ Implementations shall not allow `cl_kernel` objects to hold reference counts to `cl_kernel` arguments, because no mechanism is provided for the user to tell the kernel to release that ownership right. If the kernel holds ownership rights on kernel args, that would make it impossible for the user to tell with certainty when he may safely release user allocated resources associated with OpenCL objects such as the `cl_mem` backing store used with `CL_MEM_USE_HOST_PTR`.

- ✚ CL_INVALID_MEM_OBJECT for an argument declared to be a memory object when the specified *arg_value* is not a valid memory object.
- ✚ CL_INVALID_SAMPLER for an argument declared to be of type *sampler_t* when the specified *arg_value* is not a valid sampler object.
- ✚ CL_INVALID_ARG_SIZE if *arg_size* does not match the size of the data type for an argument that is not a memory object or if the argument is a memory object and *arg_size* \neq `sizeof(cl_mem)` or if *arg_size* is zero and the argument is declared with the *__local* qualifier or if the argument is a sampler and *arg_size* \neq `sizeof(cl_sampler)`.
- ✚ CL_INVALID_ARG_VALUE if the argument is an image declared with the *read_only* qualifier and *arg_value* refers to an image object created with *cl_mem_flags* of CL_MEM_WRITE or if the image argument is declared with the *write_only* qualifier and *arg_value* refers to an image object created with *cl_mem_flags* of CL_MEM_READ.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.7.3 Kernel Object Queries

The function

```
cl_int          clGetKernelInfo (cl_kernel kernel,
                                cl_kernel_info param_name,
                                size_t param_value_size,
                                void *param_value,
                                size_t *param_value_size_ret)
```

returns information about the kernel object.

kernel specifies the kernel object being queried.

param_name specifies the information to query. The list of supported *param_name* types and the information returned in *param_value* by **clGetKernelInfo** is described in *table 5.15*.

param_value is a pointer to memory where the appropriate result being queried is returned. If *param_value* is NULL, it is ignored.

param_value_size is used to specify the size in bytes of memory pointed to by *param_value*. This size must be \geq size of return type as described in *table 5.15*.

param_value_size_ret returns the actual size in bytes of data copied to *param_value*. If *param_value_size_ret* is NULL, it is ignored.

| cl kernel info | Return Type | Info. returned in <i>param_value</i> |
|---|--------------------|---|
| CL_KERNEL_FUNCTION_NAME | char[] | Return the kernel function name. |
| CL_KERNEL_NUM_ARGS | cl_uint | Return the number of arguments to <i>kernel</i> . |
| CL_KERNEL_REFERENCE_COUNT¹⁷ | cl_uint | Return the <i>kernel</i> reference count. |
| CL_KERNEL_CONTEXT | cl_context | Return the context associated with <i>kernel</i> . |
| CL_KERNEL_PROGRAM | cl_program | Return the program object associated with <i>kernel</i> . |
| CL_KERNEL_ATTRIBUTES | char[] | <p>Returns any attributes specified using the <code>__attribute__</code> qualifier with the kernel function declaration in the program source. These attributes include attributes described in <i>section 6.11.2</i> and other attributes supported by an implementation.</p> <p>Attributes are returned as they were declared inside <code>__attribute__((...))</code>, with any surrounding whitespace and embedded newlines removed. When multiple attributes are present, they are returned as a single, space delimited string.</p> |

Table 5.15 *clGetKernelInfo* parameter queries.

clGetKernelInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_VALUE if *param_name* is not valid, or if size in bytes specified by *param_value_size* is < size of return type as described in *table 5.15* and *param_value* is not NULL.
- ✚ CL_INVALID_KERNEL if *kernel* is a not a valid kernel object.

¹⁷ The reference count returned should be considered immediately stale. It is unsuitable for general use in applications. This feature is provided for identifying memory leaks.

- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```
cl_int  clGetKernelWorkGroupInfo (cl_kernel kernel,
                                   cl_device_id device,
                                   cl_kernel_work_group_info param_name,
                                   size_t param_value_size,
                                   void *param_value,
                                   size_t *param_value_size_ret)
```

returns information about the kernel object that may be specific to a device.

kernel specifies the kernel object being queried.

device identifies a specific device in the list of devices associated with *kernel*. The list of devices is the list of devices in the OpenCL context that is associated with *kernel*. If the list of devices associated with *kernel* is a single device, *device* can be a NULL value.

param_name specifies the information to query. The list of supported *param_name* types and the information returned in *param_value* by **clGetKernelWorkGroupInfo** is described in *table 5.16*.

param_value is a pointer to memory where the appropriate result being queried is returned. If *param_value* is NULL, it is ignored.

param_value_size is used to specify the size in bytes of memory pointed to by *param_value*. This size must be \geq size of return type as described in *table 5.16*.

param_value_size_ret returns the actual size in bytes of data copied to *param_value*. If *param_value_size_ret* is NULL, it is ignored.

| cl kernel work group info | Return Type | Info. returned in <i>param_value</i> |
|----------------------------|-------------|---|
| CL_KERNEL_GLOBAL_WORK_SIZE | size_t[3] | This provides a mechanism for the application to query the maximum global size that can be used to execute a kernel (i.e. <i>global_work_size</i> argument to clEnqueueNDRangeKernel) on a custom device given by <i>device</i> or a built-in kernel on an OpenCL device given by <i>device</i> . |

| | | |
|---|-----------|--|
| | | If <i>device</i> is not a custom device or <i>kernel</i> is not a built-in kernel, clGetKernelArgInfo returns the error CL_INVALID_VALUE. |
| CL_KERNEL_WORK_GROUP_SIZE | size_t | This provides a mechanism for the application to query the maximum work-group size that can be used to execute a kernel on a specific device given by <i>device</i> . The OpenCL implementation uses the resource requirements of the kernel (register usage etc.) to determine what this work-group size should be. |
| CL_KERNEL_COMPILE_WORK_GROUP_SIZE | size_t[3] | Returns the work-group size specified by the <code>__attribute__((reqd_work_group_size(X, Y, Z)))</code> qualifier. Refer to <i>section 6.7.2</i> . If the work-group size is not specified using the above attribute qualifier (0, 0, 0) is returned. |
| CL_KERNEL_LOCAL_MEM_SIZE | cl_ulong | Returns the amount of local memory in bytes being used by a kernel. This includes local memory that may be needed by an implementation to execute the kernel, variables declared inside the kernel with the <code>__local</code> address qualifier and local memory to be allocated for arguments to the kernel declared as pointers with the <code>__local</code> address qualifier and whose size is specified with clSetKernelArg . If the local memory size, for any pointer argument to the kernel declared with the <code>__local</code> address qualifier, is not specified, its size is assumed to be 0. |
| CL_KERNEL_PREFERRED_WORK_GROUP_SIZE_MULTIPLE | size_t | Returns the preferred multiple of work-group size for launch. This is a performance hint. Specifying a work-group size that is not a multiple of the value returned by this query as the value |

| | | |
|-----------------------------------|----------|---|
| | | of the local work size argument to clEnqueueNDRangeKernel will not fail to enqueue the kernel for execution unless the work-group size specified is larger than the device maximum. |
| CL_KERNEL_PRIVATE_MEM_SIZE | cl_ulong | Returns the minimum amount of private memory, in bytes, used by each work-item in the kernel. This value may include any private memory needed by an implementation to execute the kernel, including that used by the language built-ins and variable declared inside the kernel with the <code>__private</code> qualifier. |

Table 5.16 *clGetKernelWorkGroupInfo* parameter queries.

clGetKernelWorkGroupInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_DEVICE if *device* is not in the list of devices associated with *kernel* or if *device* is NULL but there is more than one device associated with *kernel*.
- ✚ CL_INVALID_VALUE if *param_name* is not valid, or if size in bytes specified by *param_value_size* is < size of return type as described in *table 5.15* and *param_value* is not NULL.
- ✚ CL_INVALID_VALUE if *param_name* is CL_KERNEL_GLOBAL_WORK_SIZE and *device* is not a custom device or *kernel* is not a built-in kernel.
- ✚ CL_INVALID_KERNEL if *kernel* is a not a valid kernel object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```

cl_int clGetKernelArgInfo (cl_kernel kernel,
                           cl_uint arg_indx,
                           cl_kernel_arg_info param_name,
                           size_t param_value_size,
                           void *param_value,
                           size_t *param_value_size_ret)

```

returns information about the arguments of a kernel. Kernel argument information is only available if the program object associated with *kernel* is created with **clCreateProgramWithSource** and the program executable is built with the **-cl-kernel-arg-info** option specified in *options* argument to **clBuildProgram** or **clCompileProgram**.

kernel specifies the kernel object being queried.

arg_indx is the argument index. Arguments to the kernel are referred by indices that go from 0 for the leftmost argument to *n* - 1, where *n* is the total number of arguments declared by a kernel.

param_name specifies the argument information to query. The list of supported *param_name* types and the information returned in *param_value* by **clGetKernelArgInfo** is described in [table 5.17](#).

param_value is a pointer to memory where the appropriate result being queried is returned. If *param_value* is NULL, it is ignored.

param_value_size is used to specify the size in bytes of memory pointed to by *param_value*. This size must be > size of return type as described in [table 5.17](#). *param_value_size_ret* returns the actual size in bytes of data copied to *param_value*. If *param_value_size_ret* is NULL, it is ignored.

| cl_kernel_arg_info | Return Type | Info. returned in <i>param_value</i> |
|--|---------------------------------|---|
| CL_KERNEL_ARG_ADDRESS_QUALIFIER | cl_kernel_arg_address_qualifier | Returns the address qualifier specified for the argument given by <i>arg_indx</i> . This can be one of the following values: CL_KERNEL_ARG_ADDRESS_GLOBAL CL_KERNEL_ARG_ADDRESS_LOCAL CL_KERNEL_ARG_ADDRESS_CONSTANT CL_KERNEL_ARG_ADDRESS_PRIVATE If no address qualifier is specified, the default address qualifier which is CL_KERNEL_ARG_ADDRESS_PRIVATE is returned. |
| CL_KERNEL_ARG_ACCESS_QUALIFIER | cl_kernel_arg_access_qualifier | Returns the access qualifier specified for the argument given by <i>arg_indx</i> . This can be one of the following values: CL_KERNEL_ARG_ACCESS_READ_ONLY CL_KERNEL_ARG_ACCESS_WRITE_ONLY CL_KERNEL_ARG_ACCESS_READ_WRITE CL_KERNEL_ARG_ACCESS_NONE If argument is not an image type, |

| | | |
|-------------------------------------|------------------------------|--|
| | | CL_KERNEL_ARG_ACCESS_NONE is returned. If argument is an image type, the access qualifier specified or the default access qualifier is returned. |
| CL_KERNEL_ARG_TYPE_NAME | char[] | Returns the type name specified for the argument given by <i>arg_indx</i> . The type name returned will be the argument type name as it was declared with any whitespace removed. If argument type name is an unsigned scalar type (i.e. unsigned char, unsigned short, unsigned int, unsigned long), uchar, ushort, uint and ulong will be returned. The argument type name returned does not include any type qualifiers. |
| CL_KERNEL_ARG_TYPE_QUALIFIER | cl_kernel_arg_type_qualifier | <p>Returns the type qualifier specified for the argument given by <i>arg_indx</i>. The returned value can be:</p> <p>CL_KERNEL_ARG_TYPE_CONST CL_KERNEL_ARG_TYPE_RESTRICT CL_KERNEL_ARG_TYPE_VOLATILE, a combination of the above enums or CL_KERNEL_ARG_TYPE_NONE</p> <p>NOTE: CL_KERNEL_ARG_TYPE_VOLATILE is returned if the argument is a pointer and the referenced type is declared with the volatile qualifier. For example, a kernel argument declared as <code>global int volatile *x</code> returns CL_KERNEL_ARG_TYPE_VOLATILE but a kernel argument declared as <code>global int *volatile x</code> does not. Similarly, CL_KERNEL_ARG_TYPE_RESTRICT or CL_KERNEL_ARG_TYPE_CONST is returned if the argument is a pointer and the referenced type is declared with the restrict or const qualifier. For example, a kernel argument declared as <code>global int const *x</code> returns CL_KERNEL_ARG_TYPE_CONST but a kernel argument declared as <code>global int *const x</code> does not.</p> <p>If the argument is declared with the constant address space qualifier, the CL_KERNEL_ARG_TYPE_CONST type qualifier will be set.</p> |
| CL_KERNEL_ARG_NAME | char[] | Returns the name specified for the argument given by <i>arg_indx</i> . |

Table 5.17 *clGetKernelArgInfo* parameter queries.

clGetKernelArgInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_ARG_INDEX if *arg_indx* is not a valid argument index.
- ✚ CL_INVALID_VALUE if *param_name* is not valid, or if size in bytes specified by *param_value* size is < size of return type as described in *table 5.17* and *param_value* is not NULL.
- ✚ CL_KERNEL_ARG_INFO_NOT_AVAILABLE if the argument information is not available for kernel.
- ✚ CL_INVALID_KERNEL if *kernel* is a not a valid kernel object.

5.8 Executing Kernels

The function

```
cl_int      clEnqueueNDRangeKernel (cl_command_queue command_queue,
                                     cl_kernel kernel,
                                     cl_uint work_dim,
                                     const size_t *global_work_offset,
                                     const size_t *global_work_size,
                                     const size_t *local_work_size,
                                     cl_uint num_events_in_wait_list,
                                     const cl_event *event_wait_list,
                                     cl_event *event)
```

enqueues a command to execute a kernel on a device.

command_queue is a valid command-queue. The kernel will be queued for execution on the device associated with *command_queue*.

kernel is a valid kernel object. The OpenCL context associated with *kernel* and *command_queue* must be the same.

work_dim is the number of dimensions used to specify the global work-items and work-items in the work-group. *work_dim* must be greater than zero and less than or equal to CL_DEVICE_MAX_WORK_ITEM_DIMENSIONS.

global_work_offset can be used to specify an array of *work_dim* unsigned values that describe the offset used to calculate the global ID of a work-item. If *global_work_offset* is NULL, the global IDs start at offset (0, 0, ... 0).

global_work_size points to an array of *work_dim* unsigned values that describe the number of global work-items in *work_dim* dimensions that will execute the kernel function. The total number of global work-items is computed as *global_work_size*[0] * ... * *global_work_size*[*work_dim* - 1].

local_work_size points to an array of *work_dim* unsigned values that describe the number of work-items that make up a work-group (also referred to as the size of the work-group) that will execute the kernel specified by *kernel*. The total number of work-items in a work-group is computed as *local_work_size*[0] * ... * *local_work_size*[*work_dim* - 1]. The total number of work-items in the work-group must be less than or equal to the CL_DEVICE_MAX_WORK_GROUP_SIZE value specified in *table 4.3* and the number of work-items specified in *local_work_size*[0], ... *local_work_size*[*work_dim* - 1] must be less than or equal to the corresponding values specified by CL_DEVICE_MAX_WORK_ITEM_SIZES[0], ..., CL_DEVICE_MAX_WORK_ITEM_SIZES[*work_dim* - 1]. The explicitly specified *local_work_size* will be used to determine how to break the global work-items specified by

global_work_size into appropriate work-group instances. If *local_work_size* is specified, the values specified in *global_work_size*[0], ... *global_work_size*[*work_dim* - 1] must be evenly divisible by the corresponding values specified in *local_work_size*[0], ... *local_work_size*[*work_dim* - 1].

The work-group size to be used for *kernel* can also be specified in the program source using the `__attribute__((reqd_work_group_size(X, Y, Z)))` qualifier (refer to *section 6.7.2*). In this case the size of work group specified by *local_work_size* must match the value specified by the `reqd_work_group_size` attribute qualifier.

local_work_size can also be a NULL value in which case the OpenCL implementation will determine how to be break the global work-items into appropriate work-group instances.

These work-group instances are executed in parallel across multiple compute units or concurrently on the same compute unit.

Each work-item is uniquely identified by a global identifier. The global ID, which can be read inside the kernel, is computed using the value given by *global_work_size* and *global_work_offset*. In addition, a work-item is also identified within a work-group by a unique local ID. The local ID, which can also be read by the kernel, is computed using the value given by *local_work_size*. The starting local ID is always (0, 0, ... 0).

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular kernel execution instance. Event objects are unique and can be used to identify a particular kernel execution instance later on. If *event* is NULL, no event will be created for this kernel execution instance and therefore it will not be possible for the application to query or queue a wait for this particular kernel execution instance. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueNDRangeKernel returns CL_SUCCESS if the kernel execution was successfully queued. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_PROGRAM_EXECUTABLE if there is no successfully built program executable available for device associated with *command_queue*.
- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.

- ✚ CL_INVALID_KERNEL if *kernel* is not a valid kernel object.
- ✚ CL_INVALID_CONTEXT if context associated with *command_queue* and *kernel* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_KERNEL_ARGS if the kernel argument values have not been specified.
- ✚ CL_INVALID_WORK_DIMENSION if *work_dim* is not a valid value (i.e. a value between 1 and 3).
- ✚ CL_INVALID_GLOBAL_WORK_SIZE if *global_work_size* is NULL, or if any of the values specified in *global_work_size*[0], ... *global_work_size*[*work_dim* - 1] are 0 or exceed the range given by the `sizeof(size_t)` for the device on which the kernel execution will be enqueued.
- ✚ CL_INVALID_GLOBAL_OFFSET if the value specified in *global_work_size* + the corresponding values in *global_work_offset* for any dimensions is greater than the `sizeof(size_t)` for the device on which the kernel execution will be enqueued.
- ✚ CL_INVALID_WORK_GROUP_SIZE if *local_work_size* is specified and number of work-items specified by *global_work_size* is not evenly divisible by size of work-group given by *local_work_size* or does not match the work-group size specified for *kernel* using the `__attribute__((reqd_work_group_size(X, Y, Z)))` qualifier in program source.
- ✚ CL_INVALID_WORK_GROUP_SIZE if *local_work_size* is specified and the total number of work-items in the work-group computed as *local_work_size*[0] * ... *local_work_size*[*work_dim* - 1] is greater than the value specified by CL_DEVICE_MAX_WORK_GROUP_SIZE in *table 4.3*.
- ✚ CL_INVALID_WORK_GROUP_SIZE if *local_work_size* is NULL and the `__attribute__((reqd_work_group_size(X, Y, Z)))` qualifier is used to declare the work-group size for *kernel* in the program source.
- ✚ CL_INVALID_WORK_ITEM_SIZE if the number of work-items specified in any of *local_work_size*[0], ... *local_work_size*[*work_dim* - 1] is greater than the corresponding values specified by CL_DEVICE_MAX_WORK_ITEM_SIZES[0], CL_DEVICE_MAX_WORK_ITEM_SIZES[*work_dim* - 1].
- ✚ CL_MISALIGNED_SUB_BUFFER_OFFSET if a sub-buffer object is specified as the value for an argument that is a buffer object and the *offset* specified when the sub-buffer object is created is not aligned to CL_DEVICE_MEM_BASE_ADDR_ALIGN value for device associated with *queue*.

- ✚ CL_INVALID_IMAGE_SIZE if an image object is specified as an argument value and the image dimensions (image width, height, specified or compute row and/or slice pitch) are not supported by device associated with *queue*.
- ✚ CL_IMAGE_FORMAT_NOT_SUPPORTED if an image object is specified as an argument value and the image format (image channel order and data type) is not supported by device associated with *queue*.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to queue the execution instance of *kernel* on the command-queue because of insufficient resources needed to execute the kernel. For example, the explicitly specified *local_work_size* causes a failure to execute the kernel because of insufficient resources such as registers or local memory. Another example would be the number of read-only image args used in *kernel* exceed the CL_DEVICE_MAX_READ_IMAGE_ARGS value for device or the number of write-only image args used in *kernel* exceed the CL_DEVICE_MAX_WRITE_IMAGE_ARGS value for device or the number of samplers used in *kernel* exceed CL_DEVICE_MAX_SAMPLERS for device.
- ✚ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with image or buffer objects specified as arguments to *kernel*.
- ✚ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```

cl_int      clEnqueueTask (cl_command_queue command_queue,
                          cl_kernel kernel,
                          cl_uint num_events_in_wait_list,
                          const cl_event *event_wait_list,
                          cl_event *event)

```

enqueues a command to execute a kernel on a device. The kernel is executed using a single work-item.

command_queue is a valid command-queue. The kernel will be queued for execution on the device associated with *command_queue*.

kernel is a valid kernel object. The OpenCL context associated with *kernel* and *command-queue* must be the same.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed. If *event_wait_list* is NULL, then this particular command does not wait on any event to complete. If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

event returns an event object that identifies this particular kernel execution instance. Event objects are unique and can be used to identify a particular kernel execution instance later on. If *event* is NULL, no event will be created for this kernel execution instance and therefore it will not be possible for the application to query or queue a wait for this particular kernel execution instance. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueTask is equivalent to calling **clEnqueueNDRangeKernel** with *work_dim* = 1, *global_work_offset* = NULL, *global_work_size*[0] set to 1 and *local_work_size*[0] set to 1.

clEnqueueTask returns CL_SUCCESS if the kernel execution was successfully queued. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_PROGRAM_EXECUTABLE if there is no successfully built program executable available for device associated with *command_queue*.
- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_KERNEL if *kernel* is not a valid kernel object.
- ✚ CL_INVALID_CONTEXT if context associated with *command_queue* and *kernel* are not the same or if the context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_KERNEL_ARGS if the kernel argument values have not been specified.
- ✚ CL_INVALID_WORK_GROUP_SIZE if a work-group size is specified for *kernel* using the `__attribute__((reqd_work_group_size(X, Y, Z)))` qualifier in program source and is not (1, 1, 1).
- ✚ CL_MISALIGNED_SUB_BUFFER_OFFSET if a sub-buffer object is specified as the value for an argument that is a buffer object and the *offset* specified when the sub-buffer object is created is not aligned to CL_DEVICE_MEM_BASE_ADDR_ALIGN value for device associated with *queue*.

- ✚ CL_INVALID_IMAGE_SIZE if an image object is specified as an argument value and the image dimensions (image width, height, specified or compute row and/or slice pitch) are not supported by device associated with *queue*.
- ✚ CL_IMAGE_FORMAT_NOT_SUPPORTED if an image object is specified as an argument value and the image format (image channel order and data type) is not supported by device associated with *queue*.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to queue the execution instance of *kernel* on the command-queue because of insufficient resources needed to execute the kernel.
- ✚ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with image or buffer objects specified as arguments to *kernel*.
- ✚ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```

cl_int  clEnqueueNativeKernel (cl_command_queue command_queue,
                                void (CL_CALLBACK *user_func)(void *)
                                void *args,
                                size_t cb_args,
                                cl_uint num_mem_objects,
                                const cl_mem *mem_list,
                                const void **args_mem_loc,
                                cl_uint num_events_in_wait_list,
                                const cl_event *event_wait_list,
                                cl_event *event)

```

enqueues a command to execute a native C/C++ function not compiled using the OpenCL compiler.

command_queue is a valid command-queue. A native user function can only be executed on a command-queue created on a device that has CL_EXEC_NATIVE_KERNEL capability set in CL_DEVICE_EXECUTION_CAPABILITIES as specified in *table 4.3*.

user_func is a pointer to a host-callable user function.

args is a pointer to the args list that *user_func* should be called with.

cb_args is the size in bytes of the args list that *args* points to.

The data pointed to by *args* and *cb_args* bytes in size will be copied and a pointer to this copied region will be passed to *user_func*. The copy needs to be done because the memory objects (*cl_mem* values) that *args* may contain need to be modified and replaced by appropriate pointers to global memory. When **clEnqueueNativeKernel** returns, the memory region pointed to by *args* can be reused by the application.

num_mem_objects is the number of buffer objects that are passed in *args*.

mem_list is a list of valid buffer objects, if *num_mem_objects* > 0. The buffer object values specified in *mem_list* are memory object handles (*cl_mem* values) returned by **clCreateBuffer** or NULL.

args_mem_loc is a pointer to appropriate locations that *args* points to where memory object handles (*cl_mem* values) are stored. Before the user function is executed, the memory object handles are replaced by pointers to global memory.

event_wait_list, *num_events_in_wait_list* and *event* are as described in **clEnqueueNDRangeKernel**.

clEnqueueNativeKernel returns CL_SUCCESS if the user function execution instance was successfully queued. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_CONTEXT if context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_VALUE if *user_func* is NULL.
- ✚ CL_INVALID_VALUE if *args* is a NULL value and *cb_args* > 0, or if *args* is a NULL value and *num_mem_objects* > 0.
- ✚ CL_INVALID_VALUE if *args* is not NULL and *cb_args* is 0.
- ✚ CL_INVALID_VALUE if *num_mem_objects* > 0 and *mem_list* or *args_mem_loc* are NULL.
- ✚ CL_INVALID_VALUE if *num_mem_objects* = 0 and *mem_list* or *args_mem_loc* are not NULL.

- ✚ CL_INVALID_OPERATION if the device associated with *command_queue* cannot execute the native kernel.
- ✚ CL_INVALID_MEM_OBJECT if one or more memory objects specified in *mem_list* are not valid or are not buffer objects.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to queue the execution instance of *kernel* on the command-queue because of insufficient resources needed to execute the kernel.
- ✚ CL_MEM_OBJECT_ALLOCATION_FAILURE if there is a failure to allocate memory for data store associated with buffer objects specified as arguments to *kernel*.
- ✚ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

NOTE:

The total number of read-only images specified as arguments to a kernel cannot exceed CL_DEVICE_MAX_READ_IMAGE_ARGS. Each 2D image array argument to a kernel declared with the *read_only* qualifier counts as one image.

The total number of write-only images specified as arguments to a kernel cannot exceed CL_DEVICE_MAX_WRITE_IMAGE_ARGS. Each 2D image array argument to a kernel declared with the *write_only* qualifier counts as one image.

5.9 Event Objects

Event objects can be used to refer to a kernel execution command (**clEnqueueNDRangeKernel**, **clEnqueueTask**, **clEnqueueNativeKernel**), read, write, map and copy commands on memory objects (**clEnqueue{Read|Write|Map}Buffer**, **clEnqueueUnmapMemObject**, **clEnqueue{Read|Write}BufferRect**, **clEnqueue{Read|Write|Map}Image**, **clEnqueueCopy{Buffer|Image}**, **clEnqueueCopyBufferRect**, **clEnqueueCopyBufferToImage**, **clEnqueueCopyImageToBuffer**), **clEnqueueMarkerWithWaitList**, **clEnqueueBarrierWithWaitList** (refer to *section 5.10*) or user events.

An event object can be used to track the execution status of a command. The API calls that enqueue commands to a command-queue create a new event object that is returned in the *event* argument. In case of an error enqueueing the command in the command-queue the event argument does not return an event object.

The execution status of an enqueued command at any given point in time can be one of the following:

- ✚ CL_QUEUED – This indicates that the command has been enqueued in a command-queue. This is the initial state of all events except user events.
- ✚ CL_SUBMITTED – This is the initial state for all user events. For all other events, this indicates that the command has been submitted by the host to the device.
- ✚ CL_RUNNING – This indicates that the device has started executing this command. In order for the execution status of an enqueued command to change from CL_SUBMITTED to CL_RUNNING, all events that this command is waiting on must have completed successfully i.e. their execution status must be CL_COMPLETE.
- ✚ CL_COMPLETE – This indicates that the command has successfully completed.
- ✚ Error code – The error code is a negative integer value and indicates that the command was abnormally terminated. Abnormal termination may occur for a number of reasons such as a bad memory access.

NOTE: A command is considered to be complete if its execution status is CL_COMPLETE or is a negative integer value.

If the execution of a command is terminated, the command-queue associated with this terminated command, and the associated context (and all other command-queues in this context) may no longer be available. The behavior of OpenCL API calls that use this context (and command-queues associated with this context) are now considered to be implementation-defined. The user registered callback function specified when context is created can be used to report appropriate error information.

The function

```
cl_event      clCreateUserEvent (cl_context context, cl_int *errcode_ret)
```

creates a user event object. User events allow applications to enqueue commands that wait on a user event to finish before the command is executed by the device.

context must be a valid OpenCL context.

errcode_ret will return an appropriate error code. If *errcode_ret* is NULL, no error code is returned.

clCreateUserEvent returns a valid non-zero event object and *errcode_ret* is set to CL_SUCCESS if the user event object is created successfully. Otherwise, it returns a NULL value with one of the following error values returned in *errcode_ret*:

- ✚ CL_INVALID_CONTEXT if *context* is not a valid context.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The execution status of the user event object created is set to CL_SUBMITTED.

The function

```
cl_int      clSetUserEventStatus (cl_event event, cl_int execution_status)
```

sets the execution status of a user event object.

event is a user event object created using **clCreateUserEvent**.

execution_status specifies the new execution status to be set and can be CL_COMPLETE or a negative integer value to indicate an error. A negative integer value causes all enqueued commands that wait on this user event to be terminated. **clSetUserEventStatus** can only be called once to change the execution status of *event*.

clSetUserEventStatus returns CL_SUCCESS if the function was executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_EVENT if *event* is not a valid user event object.

- ✚ CL_INVALID_VALUE if the *execution_status* is not CL_COMPLETE or a negative integer value.
- ✚ CL_INVALID_OPERATION if the *execution_status* for *event* has already been changed by a previous call to **clSetUserEventStatus**.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

NOTE: Enqueued commands that specify user events in the *event_wait_list* argument of **clEnqueue***** commands must ensure that the status of these user events being waited on are set using **clSetUserEventStatus** before any OpenCL APIs that release OpenCL objects except for event objects are called; otherwise the behavior is undefined.

For example, the following code sequence will result in undefined behavior of **clReleaseMemObject**.

```
ev1 = clCreateUserEvent(ctx, NULL);
clEnqueueWriteBuffer(cq, buf1, CL_FALSE, ...,
                    1, &ev1, NULL);
clEnqueueWriteBuffer(cq, buf2, CL_FALSE, ...);
clReleaseMemObject(buf2);
clSetUserEventStatus(ev1, CL_COMPLETE);
```

The following code sequence, however, works correctly.

```
ev1 = clCreateUserEvent(ctx, NULL);
clEnqueueWriteBuffer(cq, buf1, CL_FALSE, ...,
                    1, &ev1, NULL);
clEnqueueWriteBuffer(cq, buf2, CL_FALSE, ...);
clSetUserEventStatus(ev1, CL_COMPLETE);
clReleaseMemObject(buf2);
```

The function

```
cl_int          clWaitForEvents (cl_uint num_events, const cl_event *event_list)
```

waits on the host thread for commands identified by event objects in *event_list* to complete. A command is considered complete if its execution status is CL_COMPLETE or a negative value. The events specified in *event_list* act as synchronization points.

clWaitForEvents returns CL_SUCCESS if the execution status of all events in *event_list* is CL_COMPLETE. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_VALUE if *num_events* is zero or *event_list* is NULL.
- ✚ CL_INVALID_CONTEXT if events specified in *event_list* do not belong to the same context.
- ✚ CL_INVALID_EVENT if event objects specified in *event_list* are not valid event objects.
- ✚ CL_EXEC_STATUS_ERROR_FOR_EVENTS_IN_WAIT_LIST if the execution status of any of the events in *event_list* is a negative integer value.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```

cl_int      clGetEventInfo (cl_event event,
                             cl_event_info param_name,
                             size_t param_value_size,
                             void *param_value,
                             size_t *param_value_size_ret)

```

returns information about the event object.

event specifies the event object being queried.

param_name specifies the information to query. The list of supported *param_name* types and the information returned in *param_value* by **clGetEventInfo** is described in *table 5.18*.

param_value is a pointer to memory where the appropriate result being queried is returned. If *param_value* is NULL, it is ignored.

param_value_size is used to specify the size in bytes of memory pointed to by *param_value*. This size must be \geq size of return type as described in *table 5.18*.

param_value_size_ret returns the actual size in bytes of data copied to *param_value*. If *param_value_size_ret* is NULL, it is ignored.

| cl_event_info | Return Type | Info. returned in <i>param_value</i> |
|-------------------|-------------|--|
| CL_EVENT_COMMAND_ | cl_command_ | Return the command-queue associated with |

| | | |
|---|-----------------|--|
| QUEUE | queue | <i>event</i> . For user event objects, a NULL value is returned. |
| CL_EVENT_CONTEXT | cl_context | Return the context associated with <i>event</i> . |
| CL_EVENT_COMMAND_TYPE | cl_command_type | Return the command associated with <i>event</i> . Can be one of the following values: CL_COMMAND_NDRANGE_KERNEL CL_COMMAND_TASK CL_COMMAND_NATIVE_KERNEL CL_COMMAND_READ_BUFFER CL_COMMAND_WRITE_BUFFER CL_COMMAND_COPY_BUFFER CL_COMMAND_READ_IMAGE CL_COMMAND_WRITE_IMAGE CL_COMMAND_COPY_IMAGE CL_COMMAND_COPY_BUFFER_TO_IMAGE CL_COMMAND_COPY_IMAGE_TO_BUFFER CL_COMMAND_MAP_BUFFER CL_COMMAND_MAP_IMAGE CL_COMMAND_UNMAP_MEM_OBJECT CL_COMMAND_MARKER CL_COMMAND_ACQUIRE_GL_OBJECTS CL_COMMAND_RELEASE_GL_OBJECTS CL_COMMAND_READ_BUFFER_RECT CL_COMMAND_WRITE_BUFFER_RECT CL_COMMAND_COPY_BUFFER_RECT CL_COMMAND_USER CL_COMMAND_BARRIER CL_COMMAND_MIGRATE_MEM_OBJECTS CL_COMMAND_FILL_BUFFER CL_COMMAND_FILL_IMAGE |
| CL_EVENT_COMMAND_EXECUTION_STATUS¹⁸ | cl_int | Return the execution status of the command identified by <i>event</i> . Valid values are: CL_QUEUED (command has been enqueued in the command-queue), CL_SUBMITTED (enqueued command has been submitted by the host to the device associated with the command-queue), CL_RUNNING (device is currently executing this command), |

¹⁸ The error code values are negative, and event state values are positive. The event state values are ordered from the largest value (CL_QUEUED) for the first or initial state to the smallest value (CL_COMPLETE or negative integer value) for the last or complete state. The value of CL_COMPLETE and CL_SUCCESS are the same.

| | | |
|--|---------|--|
| | | CL_COMPLETE (the command has completed), or Error code given by a negative integer value. (command was abnormally terminated – this may be caused by a bad memory access etc.). These error codes come from the same set of error codes that are returned from the platform or runtime API calls as return values or <i>errcode_ret</i> values. |
| CL_EVENT_REFERENCE_COUNT ¹⁹ | cl_uint | Return the <i>event</i> reference count. |

Table 5.18 *clGetEventInfo* parameter queries.

Using **clGetEventInfo** to determine if a command identified by *event* has finished execution (i.e. CL_EVENT_COMMAND_EXECUTION_STATUS returns CL_COMPLETE) is not a synchronization point. There are no guarantees that the memory objects being modified by command associated with *event* will be visible to other enqueued commands.

clGetEventInfo returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_VALUE if *param_name* is not valid, or if size in bytes specified by *param_value_size* is < size of return type as described in *table 5.18* and *param_value* is not NULL.
- ✚ CL_INVALID_VALUE if information to query given in *param_name* cannot be queried for *event*.
- ✚ CL_INVALID_EVENT if *event* is a not a valid event object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

¹⁹ The reference count returned should be considered immediately stale. It is unsuitable for general use in applications. This feature is provided for identifying memory leaks.

The function

```
cl_int      clSetEventCallback (cl_event event,
                                cl_int command_exec_callback_type,
                                void (CL_CALLBACK *pfn_event_notify)(cl_event event,
                                cl_int event_command_exec_status,
                                void *user_data),
                                void *user_data)
```

registers a user callback function for a specific command execution status. The registered callback function will be called when the execution status of command associated with *event* changes to an execution status equal to or past the status specified by *command_exec_status*.

Each call to **clSetEventCallback** registers the specified user callback function on a callback stack associated with *event*. The order in which the registered user callback functions are called is undefined.

event is a valid event object.

command_exec_callback_type specifies the command execution status for which the callback is registered. The command execution callback values for which a callback can be registered are: CL_SUBMITTED, CL_RUNNING or CL_COMPLETE²⁰. There is no guarantee that the callback functions registered for various execution status values for an event will be called in the exact order that the execution status of a command changes. Furthermore, it should be noted that receiving a call back for an event with a status other than CL_COMPLETE, in no way implies that the memory model or execution model as defined by the OpenCL specification has changed. For example, it is not valid to assume that a corresponding memory transfer has completed unless the event is in a state CL_COMPLETE.

pfn_event_notify is the event callback function that can be registered by the application. This callback function may be called asynchronously by the OpenCL implementation. It is the application's responsibility to ensure that the callback function is thread-safe. The parameters to this callback function are:

- ✚ *event* is the event object for which the callback function is invoked.
- ✚ *event_command_exec_status* represents the execution status of command for which this callback function is invoked. Refer to *table 5.18* for the command execution status values. If the callback is called as the result of the command associated with event being abnormally terminated, an appropriate error code for the error that caused the termination will be passed to *event_command_exec_status* instead.
- ✚ *user_data* is a pointer to user supplied data.

²⁰The callback function registered for a *command_exec_callback_type* value of CL_COMPLETE will be called when the command has completed successfully or is abnormally terminated.

user_data will be passed as the *user_data* argument when *pfn_notify* is called. *user_data* can be NULL.

All callbacks registered for an event object must be called. All enqueued callbacks shall be called before the event object is destroyed. Callbacks must return promptly. The behavior of calling expensive system routines, OpenCL API calls to create contexts or command-queues, or blocking OpenCL operations from the following list below, in a callback is undefined.

clFinish,
clWaitForEvents,
blocking calls to **clEnqueueReadBuffer**, **clEnqueueReadBufferRect**,
clEnqueueWriteBuffer, **clEnqueueWriteBufferRect**,
blocking calls to **clEnqueueReadImage** and **clEnqueueWriteImage**,
blocking calls to **clEnqueueMapBuffer** and **clEnqueueMapImage**,
blocking calls to **clBuildProgram**, **clCompileProgram** or **clLinkProgram**

If an application needs to wait for completion of a routine from the above list in a callback, please use the non-blocking form of the function, and assign a completion callback to it to do the remainder of your work. Note that when a callback (or other code) enqueues commands to a command-queue, the commands are not required to begin execution until the queue is flushed. In standard usage, blocking enqueue calls serve this role by implicitly flushing the queue. Since blocking calls are not permitted in callbacks, those callbacks that enqueue commands on a command queue should either call **clFlush** on the queue before returning or arrange for **clFlush** to be called later on another thread.

clSetEventCallback returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_EVENT if *event* is not a valid event object.
- ✚ CL_INVALID_VALUE if *pfn_event_notify* is NULL or if *command_exec_callback_type* is not CL_SUBMITTED, CL_RUNNING or CL_COMPLETE.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

`cl_int` **clRetainEvent** (*cl_event event*)

increments the *event* reference count. The OpenCL commands that return an event perform an implicit retain.

clRetainEvent returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_EVENT if *event* is not a valid event object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

To release an event, use the following function

```
cl_int      clReleaseEvent (cl_event event)
```

decrements the *event* reference count.

clReleaseEvent returns CL_SUCCESS if the function is executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_EVENT if *event* is not a valid event object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The event object is deleted once the reference count becomes zero, the specific command identified by this event has completed (or terminated) and there are no commands in the command-queues of a context that require a wait for this event to complete.

NOTE: Developers should be careful when releasing their last reference count on events created by **clCreateUserEvent** that have not yet been set to status of CL_COMPLETE or an error. If the user event was used in the event_wait_list argument passed to a clEnqueue*** API or another application host thread is waiting for it in **clWaitForEvents**, those commands and host threads will continue to wait for the event status to reach CL_COMPLETE or error, even after the user has released the object. Since in this scenario the developer has released his last reference count to the user event, it would be in principle no longer valid for him to change the status of the event to unblock all the other machinery. As a result the waiting tasks will wait forever, and associated events, cl_mem objects, command queues and contexts are likely to leak. In-order command queues caught up in this deadlock may cease to do any work.

5.10 Markers, Barriers and Waiting for Events

The function

```
cl_int          clEnqueueMarkerWithWaitList (cl_command_queue command_queue,
                                             cl_uint num_events_in_wait_list,
                                             const cl_event *event_wait_list,
                                             cl_event *event)
```

enqueues a marker command which waits for either a list of events to complete, or if the list is empty it waits for all commands previously enqueued in *command_queue* to complete before it completes. This command returns an *event* which can be waited on, i.e. this event can be waited on to insure that all events either in the *event_wait_list* or all previously enqueued commands, queued before this command to *command_queue*, have completed.

command_queue is a valid command-queue.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed.

If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

If *event_wait_list* is NULL, then this particular command waits until all previous enqueued commands to *command_queue* have completed.

event returns an event object that identifies this particular command. Event objects are unique and can be used to identify this marker command later on. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueMarkerWithWaitList returns CL_SUCCESS if the function is successfully executed. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_INVALID_CONTEXT if context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✚ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and

- num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

The function

```
cl_int          clEnqueueBarrierWithWaitList (cl_command_queue command_queue,
                                                cl_uint num_events_in_wait_list,
                                                const cl_event *event_wait_list,
                                                cl_event *event)
```

enqueues a barrier command which waits for either a list of events to complete, or if the list is empty it waits for all commands previously enqueued in *command_queue* to complete before it completes. This command blocks command execution, that is, any following commands enqueued after it do not execute until it completes. This command returns an *event* which can be waited on, i.e. this event can be waited on to insure that all events either in the *event_wait_list* or all previously enqueued commands, queued before this command to *command_queue*, have completed

command_queue is a valid command-queue.

event_wait_list and *num_events_in_wait_list* specify events that need to complete before this particular command can be executed.

If *event_wait_list* is NULL, *num_events_in_wait_list* must be 0. If *event_wait_list* is not NULL, the list of events pointed to by *event_wait_list* must be valid and *num_events_in_wait_list* must be greater than 0. The events specified in *event_wait_list* act as synchronization points. The context associated with events in *event_wait_list* and *command_queue* must be the same. The memory associated with *event_wait_list* can be reused or freed after the function returns.

If *event_wait_list* is NULL, then this particular command waits until all previous enqueued commands to *command_queue* have completed.

event returns an event object that identifies this particular command. Event objects are unique and can be used to identify this barrier command later on. *event* can be NULL in which case it will not be possible for the application to query the status of this command or queue a wait for this command to complete. If the *event_wait_list* and the *event* arguments are not NULL, the *event* argument should not refer to an element of the *event_wait_list* array.

clEnqueueBarrierWithWaitList returns CL_SUCCESS if the function is successfully executed. Otherwise, it returns one of the following errors:

- ✦ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✦ CL_INVALID_CONTEXT if context associated with *command_queue* and events in *event_wait_list* are not the same.
- ✦ CL_INVALID_EVENT_WAIT_LIST if *event_wait_list* is NULL and *num_events_in_wait_list* > 0, or *event_wait_list* is not NULL and *num_events_in_wait_list* is 0, or if event objects in *event_wait_list* are not valid events.
- ✦ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✦ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.11 Out-of-order Execution of Kernels and Memory Object Commands

The OpenCL functions that are submitted to a command-queue are enqueued in the order the calls are made but can be configured to execute in-order or out-of-order. The *properties* argument in **clCreateCommandQueue** can be used to specify the execution order.

If the `CL_QUEUE_OUT_OF_ORDER_EXEC_MODE_ENABLE` property of a command-queue is not set, the commands enqueued to a command-queue execute in order. For example, if an application calls **clEnqueueNDRangeKernel** to execute kernel A followed by a **clEnqueueNDRangeKernel** to execute kernel B, the application can assume that kernel A finishes first and then kernel B is executed. If the memory objects output by kernel A are inputs to kernel B then kernel B will see the correct data in memory objects produced by execution of kernel A. If the `CL_QUEUE_OUT_OF_ORDER_EXEC_MODE_ENABLE` property of a command-queue is set, then there is no guarantee that kernel A will finish before kernel B starts execution.

Applications can configure the commands enqueued to a command-queue to execute out-of-order by setting the `CL_QUEUE_OUT_OF_ORDER_EXEC_MODE_ENABLE` property of the command-queue. This can be specified when the command-queue is created. In out-of-order execution mode there is no guarantee that the enqueued commands will finish execution in the order they were queued. As there is no guarantee that kernels will be executed in order, i.e. based on when the **clEnqueueNDRangeKernel** calls are made within a command-queue, it is therefore possible that an earlier **clEnqueueNDRangeKernel** call to execute kernel A identified by event A may execute and/or finish later than a **clEnqueueNDRangeKernel** call to execute kernel B which was called by the application at a later point in time. To guarantee a specific order of execution of kernels, a wait on a particular event (in this case event A) can be used. The wait for event A can be specified in the *event_wait_list* argument to **clEnqueueNDRangeKernel** for kernel B.

In addition, a wait for events (**clEnqueueMarkerWithWaitList**) or a barrier (**clEnqueueBarrierWithWaitList**) command can be enqueued to the command-queue. The wait for events command ensures that previously enqueued commands identified by the list of events to wait for have finished before the next batch of commands is executed. The barrier command ensures that all previously enqueued commands in a command-queue have finished execution before the next batch of commands is executed.

Similarly, commands to read, write, copy or map memory objects that are enqueued after **clEnqueueNDRangeKernel**, **clEnqueueTask** or **clEnqueueNativeKernel** commands are not guaranteed to wait for kernels scheduled for execution to have completed (if the `CL_QUEUE_OUT_OF_ORDER_EXEC_MODE_ENABLE` property is set). To ensure correct ordering of commands, the event object returned by **clEnqueueNDRangeKernel**, **clEnqueueTask** or **clEnqueueNativeKernel** can be used to enqueue a wait for event or a barrier command can be enqueued that must complete before reads or writes to the memory object(s) occur.

5.12 Profiling Operations on Memory Objects and Kernels

This section describes profiling of OpenCL functions that are enqueued as commands to a command-queue. The specific functions²¹ being referred to are:

clEnqueue{Read|Write|Map}Buffer, **clEnqueue{Read|Write}BufferRect**,
clEnqueue{Read|Write|Map}Image, **clEnqueueUnmapMemObject**, **clEnqueueCopyBuffer**,
clEnqueueCopyBufferRect, **clEnqueueCopyImage**, **clEnqueueCopyImageToBuffer**,
clEnqueueCopyBufferToImage, **clEnqueueNDRangeKernel** , **clEnqueueTask** and
clEnqueueNativeKernel. These enqueued commands are identified by unique event objects.

Event objects can be used to capture profiling information that measure execution time of a command. Profiling of OpenCL commands can be enabled either by using a command-queue created with `CL_QUEUE_PROFILING_ENABLE` flag set in *properties* argument to **clCreateCommandQueue**.

If profiling is enabled, the function

```
cl_int          clGetEventProfilingInfo (cl_event event,  
                                         cl_profiling_info param_name,  
                                         size_t param_value_size,  
                                         void *param_value,  
                                         size_t *param_value_size_ret)
```

returns profiling information for the command associated with event.

event specifies the event object.

param_name specifies the profiling data to query. The list of supported *param_name* types and the information returned in *param_value* by **clGetEventProfilingInfo** is described in *table 5.19*.

param_value is a pointer to memory where the appropriate result being queried is returned. If *param_value* is NULL, it is ignored.

param_value_size is used to specify the size in bytes of memory pointed to by *param_value*. This size must be \geq size of return type as described in *table 5.19*.

param_value_size_ret returns the actual size in bytes of data copied to *param_value*. If *param_value_size_ret* is NULL, it is ignored.

²¹ **clEnqueueAcquireGLObjects** and **clEnqueueReleaseGLObjects** defined in *section 9.6.6* of the OpenCL 1.2 Extension Specification are also included.

| cl_profiling_info | Return Type | Info. returned in <i>param_value</i> |
|------------------------------------|--------------------|---|
| CL_PROFILING_COMMAND_QUEUED | cl_ulong | A 64-bit value that describes the current device time counter in nanoseconds when the command identified by <i>event</i> is enqueued in a command-queue by the host. |
| CL_PROFILING_COMMAND_SUBMIT | cl_ulong | A 64-bit value that describes the current device time counter in nanoseconds when the command identified by <i>event</i> that has been enqueued is submitted by the host to the device associated with the command-queue. |
| CL_PROFILING_COMMAND_START | cl_ulong | A 64-bit value that describes the current device time counter in nanoseconds when the command identified by <i>event</i> starts execution on the device. |
| CL_PROFILING_COMMAND_END | cl_ulong | A 64-bit value that describes the current device time counter in nanoseconds when the command identified by <i>event</i> has finished execution on the device. |

Table 5.19 *clGetEventProfilingInfo* parameter queries.

The unsigned 64-bit values returned can be used to measure the time in nano-seconds consumed by OpenCL commands.

OpenCL devices are required to correctly track time across changes in device frequency and power states. The `CL_DEVICE_PROFILING_TIMER_RESOLUTION` specifies the resolution of the timer i.e. the number of nanoseconds elapsed before the timer is incremented.

clGetEventProfilingInfo returns `CL_SUCCESS` if the function is executed successfully and the profiling information has been recorded. Otherwise, it returns one of the following errors:

- ✚ `CL_PROFILING_INFO_NOT_AVAILABLE` if the `CL_QUEUE_PROFILING_ENABLE` flag is not set for the command-queue, if the execution status of the command identified by *event* is not `CL_COMPLETE` or if *event* is a user event object.

- ✚ CL_INVALID_VALUE if *param_name* is not valid, or if size in bytes specified by *param_value_size* is < size of return type as described in *table 5.19* and *param_value* is not NULL.
- ✚ CL_INVALID_EVENT if *event* is a not a valid event object.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

5.13 Flush and Finish

The function

```
cl_int          clFlush (cl_command_queue command_queue)
```

issues all previously queued OpenCL commands in *command_queue* to the device associated with *command_queue*. **clFlush** only guarantees that all queued commands to *command_queue* will eventually be submitted to the appropriate device. There is no guarantee that they will be complete after **clFlush** returns.

clFlush returns CL_SUCCESS if the function call was executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

Any blocking commands queued in a command-queue and **clReleaseCommandQueue** perform an implicit flush of the command-queue. These blocking commands are **clEnqueueReadBuffer**, **clEnqueueReadBufferRect**, **clEnqueueReadImage**, with *blocking_read* set to CL_TRUE; **clEnqueueWriteBuffer**, **clEnqueueWriteBufferRect**, **clEnqueueWriteImage** with *blocking_write* set to CL_TRUE; **clEnqueueMapBuffer**, **clEnqueueMapImage** with *blocking_map* set to CL_TRUE; or **clWaitForEvents**.

To use event objects that refer to commands enqueued in a command-queue as event objects to wait on by commands enqueued in a different command-queue, the application must call a **clFlush** or any blocking commands that perform an implicit flush of the command-queue where the commands that refer to these event objects are enqueued.

The function

```
cl_int          clFinish (cl_command_queue command_queue)
```

blocks until all previously queued OpenCL commands in *command_queue* are issued to the associated device and have completed. **clFinish** does not return until all previously queued commands in *command_queue* have been processed and completed. **clFinish** is also a synchronization point.

clFinish returns CL_SUCCESS if the function call was executed successfully. Otherwise, it returns one of the following errors:

- ✚ CL_INVALID_COMMAND_QUEUE if *command_queue* is not a valid command-queue.
- ✚ CL_OUT_OF_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- ✚ CL_OUT_OF_HOST_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

6. The OpenCL C Programming Language

This section describes the OpenCL C programming language used to create kernels that are executed on OpenCL device(s). The OpenCL C programming language (also referred to as OpenCL C) is based on the ISO/IEC 9899:1999 C language specification (a.k.a. C99 specification) with specific extensions and restrictions. Please refer to the ISO/IEC 9899:1999 specification for a detailed description of the language grammar. This section describes modifications and restrictions to ISO/IEC 9899:1999 supported in OpenCL C.

6.1 Supported Data Types

The following data types are supported.

6.1.1 Built-in Scalar Data Types

Table 6.1 describes the list of built-in scalar data types.

| Type | Description |
|-------------------------------|---|
| bool ²² | A conditional data type which is either <i>true</i> or <i>false</i> . The value <i>true</i> expands to the integer constant 1 and the value <i>false</i> expands to the integer constant 0. |
| char | A signed two's complement 8-bit integer. |
| unsigned char, uchar | An unsigned 8-bit integer. |
| short | A signed two's complement 16-bit integer. |
| unsigned short, ushort | An unsigned 16-bit integer. |
| int | A signed two's complement 32-bit integer. |
| unsigned int, uint | An unsigned 32-bit integer. |
| long | A signed two's complement 64-bit integer. |
| unsigned long, ulong | An unsigned 64-bit integer. |
| float | A 32-bit floating-point. The float data type must conform to the IEEE 754 single precision storage format. |
| double ²³ | A 64-bit floating-point. The double data type must conform to the IEEE 754 double precision storage format. |
| half | A 16-bit floating-point. The half data type must conform to the |

²² When any scalar value is converted to **bool**, the result is 0 if the value compares equal to 0; otherwise, the result is 1.

²³ The double scalar type is an optional type that is supported if CL_DEVICE_DOUBLE_FP_CONFIG in table 4.3 for a device is not zero.

| | |
|------------------|--|
| | IEEE 754-2008 half precision storage format. |
| size_t | The unsigned integer type of the result of the sizeof operator. This is a 32-bit unsigned integer if <code>CL_DEVICE_ADDRESS_BITS</code> defined in <i>table 4.3</i> is 32-bits and is a 64-bit unsigned integer if <code>CL_DEVICE_ADDRESS_BITS</code> is 64-bits. |
| ptrdiff_t | A signed integer type that is the result of subtracting two pointers. This is a 32-bit signed integer if <code>CL_DEVICE_ADDRESS_BITS</code> defined in <i>table 4.3</i> is 32-bits and is a 64-bit signed integer if <code>CL_DEVICE_ADDRESS_BITS</code> is 64-bits. |
| intptr_t | A signed integer type with the property that any valid pointer to void can be converted to this type, then converted back to pointer to void , and the result will compare equal to the original pointer. This is a 32-bit signed integer if <code>CL_DEVICE_ADDRESS_BITS</code> defined in <i>table 4.3</i> is 32-bits and is a 64-bit signed integer if <code>CL_DEVICE_ADDRESS_BITS</code> is 64-bits. |
| uintptr_t | An unsigned integer type with the property that any valid pointer to void can be converted to this type, then converted back to pointer to void , and the result will compare equal to the original pointer. This is a 32-bit signed integer if <code>CL_DEVICE_ADDRESS_BITS</code> defined in <i>table 4.3</i> is 32-bits and is a 64-bit signed integer if <code>CL_DEVICE_ADDRESS_BITS</code> is 64-bits. |
| void | The void type comprises an empty set of values; it is an incomplete type that cannot be completed. |

Table 6.1 *Built-in Scalar Data Types*

Most built-in scalar data types are also declared as appropriate types in the OpenCL API (and header files) that can be used by an application. The following table describes the built-in scalar data type in the OpenCL C programming language and the corresponding data type available to the application:

| Type in OpenCL Language | API type for application |
|-----------------------------------|--------------------------|
| bool | n/a |
| char | cl_char |
| unsigned char, uchar | cl_uchar |
| short | cl_short |
| unsigned short, ushort | cl_ushort |
| int | cl_int |
| unsigned int, uint | cl_uint |
| long | cl_long |
| unsigned long, | cl_ulong |

| | |
|------------------|------------------|
| ulong | |
| float | cl float |
| double | cl double |
| half | cl half |
| size_t | n/a |
| ptrdiff_t | n/a |
| intptr_t | n/a |
| uintptr_t | n/a |
| void | void |

6.1.1.1 The half data type

The `half` data type must be IEEE 754-2008 compliant. `half` numbers have 1 sign bit, 5 exponent bits, and 10 mantissa bits. The interpretation of the sign, exponent and mantissa is analogous to IEEE 754 floating-point numbers. The exponent bias is 15. The `half` data type must represent finite and normal numbers, denormalized numbers, infinities and NaN. Denormalized numbers for the `half` data type which may be generated when converting a `float` to a `half` using `vstore_half` and converting a `half` to a `float` using `vload_half` cannot be flushed to zero. Conversions from `float` to `half` correctly round the mantissa to 11 bits of precision. Conversions from `half` to `float` are lossless; all `half` numbers are exactly representable as `float` values.

The `half` data type can only be used to declare a pointer to a buffer that contains half values. A few valid examples are given below:

```
void
bar (__global half *p)
{
    ...
}

__kernel void
foo (__global half *pg, __local half *pl)
{
    __global half *ptr;
    int offset;

    ptr = pg + offset;
    bar(ptr);
}
```

Below are some examples that are not valid usage of the `half` type:

```
half a;
half b[100];
```

```

half *p;
a = *p;    ← not allowed. must use vload_half function

```

Loads from a pointer to a `half` and stores to a pointer to a `half` can be performed using the `vload_half`, `vload_halfn`, `vloada_halfn` and `vstore_half`, `vstore_halfn`, `vstorea_halfn` functions respectively as described in *section 6.12.7*. The load functions read scalar or vector half values from memory and convert them to a scalar or vector float value. The store functions take a scalar or vector float value as input, convert it to a half scalar or vector value (with appropriate rounding mode) and write the half scalar or vector value to memory.

6.1.2 Built-in Vector Data Types²⁴

The `char`, unsigned `char`, `short`, unsigned `short`, `integer`, unsigned `integer`, `long`, unsigned `long`, float vector data types are supported. The vector data type is defined with the type name i.e. `char`, `uchar`, `short`, `ushort`, `int`, `uint`, `float`, `long`, `ulong` followed by a literal value `n` that defines the number of elements in the vector. Supported values of `n` are 2, 3, 4, 8, and 16 for all vector data types.

Table 6.2 describes the list of built-in vector data types.

| Type | Description |
|------------------------------|---|
| charn | A vector of <i>n</i> 8-bit signed two's complement integer values. |
| ucharn | A vector of <i>n</i> 8-bit unsigned integer values. |
| shortn | A vector of <i>n</i> 16-bit signed two's complement integer values. |
| ushortn | A vector of <i>n</i> 16-bit unsigned integer values. |
| intn | A vector of <i>n</i> 32-bit signed two's complement integer values. |
| uintn | A vector of <i>n</i> 32-bit unsigned integer values. |
| longn | A vector of <i>n</i> 64-bit signed two's complement integer values. |
| ulongn | A vector of <i>n</i> 64-bit unsigned integer values. |
| floatn | A vector of <i>n</i> 32-bit floating-point values. |
| doublen ²⁵ | A vector of <i>n</i> 64-bit floating-point values. |

Table 6.2 Built-in Vector Data Types

The built-in vector data types are also declared as appropriate types in the OpenCL API (and header files) that can be used by an application. The following table describes the built-in vector

²⁴ Built-in vector data types are supported by the OpenCL implementation even if the underlying compute device does not support any or all of the vector data types. These are to be converted by the device compiler to appropriate instructions that use underlying built-in types supported natively by the compute device. Refer to Appendix B for a description of the order of the components of a vector type in memory.

²⁵ The double vector type is an optional type that is supported if `CL_DEVICE_DOUBLE_FP_CONFIG` in table 4.3 for a device is not zero.

data type in the OpenCL C programming language and the corresponding data type available to the application:

| Type in OpenCL Language | API type for application |
|-------------------------|--------------------------|
| <i>charn</i> | <i>cl_charn</i> |
| <i>ucharn</i> | <i>cl_ucharn</i> |
| <i>shortn</i> | <i>cl_shortn</i> |
| <i>ushortn</i> | <i>cl_ushortn</i> |
| <i>intn</i> | <i>cl_intn</i> |
| <i>uintn</i> | <i>cl_uintn</i> |
| <i>longn</i> | <i>cl_longn</i> |
| <i>ulongn</i> | <i>cl_ulongn</i> |
| <i>floatn</i> | <i>cl_floatn</i> |
| <i>doublen</i> | <i>cl_doublen</i> |

6.1.3 Other Built-in Data Types

Table 6.3 describes the list of additional data types supported by OpenCL.

| Type | Description |
|--------------------------------|---|
| <i>image2d_t</i> | A 2D image. Refer to <i>section 6.12.14</i> for a detailed description of the built-in functions that use this type. |
| <i>image3d_t</i> | A 3D image. Refer to <i>section 6.12.14</i> for a detailed description of the built-in functions that use this type. |
| <i>image2d_array_t</i> | A 2D image array. Refer to <i>section 6.12.14</i> for a detailed description of the built-in functions that use this type. |
| <i>image1d_t</i> | A 1D image. Refer to <i>section 6.12.14</i> for a detailed description of the built-in functions that use this type. |
| <i>image1d_buffer_t</i> | A 1D image created from a buffer object. Refer to <i>section 6.12.14</i> for a detailed description of the built-in functions that use this type. |
| <i>image1d_array_t</i> | A 1D image array. Refer to <i>section 6.12.14</i> for a detailed description of the built-in functions that use this type. |
| <i>sampler_t</i> | A sampler type. Refer to <i>section 6.12.14</i> for a detailed description the built-in functions that use of this type. |
| <i>event_t</i> | An event. This can be used to identify async copies from global to local memory and vice-versa. Refer to <i>section 6.12.10</i> . |

Table 6.3 *Other Built-in Data Types*

NOTE: The **image2d_t**, **image3d_t**, **image2d_array_t**, **image1d_t**, **image1d_buffer_t**, **image1d_array_t** and **sampler_t** types are only defined if the device supports images i.e. **CL_DEVICE_IMAGE_SUPPORT** as described in *table 4.3* is **CL_TRUE**.

The C99 derived types (arrays, structs, unions, functions, and pointers), constructed from the built-in data types described in *sections 6.1.1, 6.1.2 and 6.1.3* are supported, with restrictions described in *section 6.9*.

6.1.4 Reserved Data Types

The data type names described in *table 6.4* are reserved and cannot be used by applications as type names. The vector data type names defined in *table 6.2*, but where *n* is any value other than 2, 3, 4, 8 and 16, are also reserved.

| Type | Description |
|--|--|
| booln | A boolean vector. |
| halfn | A 16-bit floating-point vector. |
| quad, quadn | A 128-bit floating-point scalar and vector. |
| complex half, complex halfn | A complex 16-bit floating-point scalar and vector. |
| imaginary half, imaginary halfn | An imaginary 16-bit floating-point scalar and vector. |
| complex float, complex floatn | A complex 32-bit floating-point scalar and vector. |
| imaginary float, imaginary floatn | An imaginary 32-bit floating-point scalar and vector. |
| complex double, complex doublen, | A complex 64-bit floating-point scalar and vector. |
| imaginary double, imaginary doublen | An imaginary 64-bit floating-point scalar and vector. |
| complex quad, complex quadn, | A complex 128-bit floating-point scalar and vector. |
| imaginary quad, imaginary quadn | An imaginary 128-bit floating-point scalar and vector. |
| floatnxm | An <i>n</i> x <i>m</i> matrix of single precision floating-point values stored in column-major order. |
| doublenxm | An <i>n</i> x <i>m</i> matrix of double precision floating-point values stored in column-major order. |
| long double long doublen | A floating-point scalar and vector type with at least as much precision and range as a double and no more precision and range than a quad. |
| long long, long longn | A 128-bit signed integer scalar and vector. |
| unsigned long long, | A 128-bit unsigned integer scalar and vector. |

| | |
|---------------------------------------|--|
| ulong long, ulong long<i>n</i> | |
|---------------------------------------|--|

Table 6.4 *Reserved Data Types*

6.1.5 Alignment of Types

A data item declared to be a data type in memory is always aligned to the size of the data type in bytes. For example, a float4 variable will be aligned to a 16-byte boundary, a char2 variable will be aligned to a 2-byte boundary.

For 3-component vector data types, the size of the data type is `4 * sizeof(component)`. This means that a 3-component vector data type will be aligned to a `4 * sizeof(component)` boundary. The **vload3** and **vstore3** built-in functions can be used to read and write, respectively, 3-component vector data types from an array of packed scalar data type.

A built-in data type that is not a power of two bytes in size must be aligned to the next larger power of two. This rule applies to built-in types only, not structs or unions.

The OpenCL compiler is responsible for aligning data items to the appropriate alignment as required by the data type. For arguments to a `__kernel` function declared to be a pointer to a data type, the OpenCL compiler can assume that the pointee is always appropriately aligned as required by the data type. The behavior of an unaligned load or store is undefined, except for the **vload*n***, **vload_half*n***, **vstore*n***, and **vstore_half*n*** functions defined in *section 6.12.7*. The vector load functions can read a vector from an address aligned to the element type of the vector. The vector store functions can write a vector to an address aligned to the element type of the vector.

6.1.6 Vector Literals

Vector literals can be used to create vectors from a list of scalars, vectors or a mixture thereof. A vector literal can be used either as a vector initializer or as a primary expression. A vector literal cannot be used as an L-value.

A vector literal is written as a parenthesized vector type followed by a parenthesized comma delimited list of parameters. A vector literal operates as an overloaded function. The forms of the function that are available is the set of possible argument lists for which all arguments have the same element type as the result vector, and the total number of elements is equal to the number of elements in the result vector. In addition, a form with a single scalar of the same type as the element type of the vector is available. For example, the following forms are available for float4:

```
(float4) ( float, float, float, float )
(float4) ( float2, float, float )
(float4) ( float, float2, float )
```

```

(float4)( float, float, float2 )
(float4)( float2, float2 )
(float4)( float3, float )
(float4)( float, float3 )

(float4)( float )

```

Operands are evaluated by standard rules for function evaluation, except that implicit scalar widening shall not occur. The order in which the operands are evaluated is undefined. The operands are assigned to their respective positions in the result vector as they appear in memory order. That is, the first element of the first operand is assigned to result.x, the second element of the first operand (or the first element of the second operand if the first operand was a scalar) is assigned to result.y, etc. In the case of the form that has a single scalar operand, the operand is replicated across all lanes of the vector.

Examples:

```

float4  f = (float4)(1.0f, 2.0f, 3.0f, 4.0f);

uint4   u = (uint4)(1); ← u will be (1, 1, 1, 1).

float4  f = (float4)((float2)(1.0f, 2.0f),
                    (float2)(3.0f, 4.0f));

float4  f = (float4)(1.0f, (float2)(2.0f, 3.0f), 4.0f);

float4  f = (float4)(1.0f, 2.0f); ← error

```

6.1.7 Vector Components

The components of vector data types with 1 ... 4 components can be addressed as <vector_data_type>.xyzw. Vector data types of type char2, uchar2, short2, ushort2, int2, uint2, long2, ulong2, and float2 can access .xy elements. Vector data types of type char3, uchar3, short3, ushort3, int3, uint3, long3, ulong3, and float3 can access .xyz elements. Vector data types of type char4, uchar4, short4, ushort4, int4, uint4, long4, ulong4, float4 can access .xyzw elements.

Accessing components beyond those declared for the vector type is an error so, for example:

```

float2 pos;
pos.x = 1.0f;    // is legal
pos.z = 1.0f;    // is illegal

float3 pos;
pos.z = 1.0f;    // is legal

```

```
pos.w = 1.0f;    // is illegal
```

The component selection syntax allows multiple components to be selected by appending their names after the period (.).

```
float4 c;  
  
c.xyzw = (float4)(1.0f, 2.0f, 3.0f, 4.0f);  
c.z = 1.0f;  
c.xy = (float2)(3.0f, 4.0f);  
c.xyz = (float3)(3.0f, 4.0f, 5.0f);
```

The component selection syntax also allows components to be permuted or replicated.

```
float4 pos = (float4)(1.0f, 2.0f, 3.0f, 4.0f);  
  
float4 swiz= pos.wzyx; // swiz = (4.0f, 3.0f, 2.0f, 1.0f)  
  
float4 dup = pos.xxxy; // dup = (1.0f, 1.0f, 2.0f, 2.0f)
```

The component group notation can occur on the left hand side of an expression. To form an l-value, swizzling must be applied to an l-value of vector type, contain no duplicate components, and it results in an l-value of scalar or vector type, depending on number of components specified. Each component must be a supported scalar or vector type.

```
float4 pos = (float4)(1.0f, 2.0f, 3.0f, 4.0f);  
  
pos.xw = (float2)(5.0f, 6.0f); // pos = (5.0f, 2.0f, 3.0f, 6.0f)  
pos.wx = (float2)(7.0f, 8.0f); // pos = (8.0f, 2.0f, 3.0f, 7.0f)  
pos.xyz = (float3)(3.0f, 5.0f, 9.0f); // pos = (3.0f, 5.0f, 9.0f, 4.0f)  
pos.xx = (float2)(3.0f, 4.0f); // illegal - 'x' used twice  
  
// illegal - mismatch between float2 and float4  
pos.xy = (float4)(1.0f, 2.0f, 3.0f, 4.0f);  
  
float4 a, b, c, d;  
float16 x;  
x = (float16)(a, b, c, d);  
x = (float16)(a.xxxx, b.xyz, c.xyz, d.xyz, a.yzw);  
  
// illegal - component a.xxxxxxx is not a valid vector type  
x = (float16)(a.xxxxxxx, b.xyz, c.xyz, d.xyz);
```

Elements of vector data types can also be accessed using a numeric index to refer to the appropriate element in the vector. The numeric indices that can be used are given in the table below:

| Vector Components | Numeric indices that can be used |
|-------------------|--|
| 2-component | 0, 1 |
| 3-component | 0, 1, 2 |
| 4-component | 0, 1, 2, 3 |
| 8-component | 0, 1, 2, 3, 4, 5, 6, 7 |
| 16-component | 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, a, A, b, B, c, C, d, D, e, E, f, F |

Table 6.5 *Numeric indices for built-in vector data types*

The numeric indices must be preceded by the letter `s` or `S`.

In the following example

```
float8      f;
```

`f.s0` refers to the 1st element of the `float8` variable `f` and `f.s7` refers to the 8th element of the `float8` variable `f`.

In the following example

```
float16     x;
```

`x.sa` (or `x.sA`) refers to the 11th element of the `float16` variable `x` and `x.sf` (or `x.sF`) refers to the 16th element of the `float16` variable `x`.

The numeric indices used to refer to an appropriate element in the vector cannot be intermixed with `.xyzw` notation used to access elements of a 1 .. 4 component vector.

For example

```
float4      f, a;
a = f.x12w;  // illegal use of numeric indices with .xyzw
a.xyzw = f.s0123; // valid
```

Vector data types can use the `.lo` (or `.even`) and `.hi` (or `.odd`) suffixes to get smaller vector types or to combine smaller vector types to a larger vector type. Multiple levels of `.lo` (or `.even`) and `.hi` (or `.odd`) suffixes can be used until they refer to a scalar term.

The `.lo` suffix refers to the lower half of a given vector. The `.hi` suffix refers to the upper half of a given vector.

The `.even` suffix refers to the even elements of a vector. The `.odd` suffix refers to the odd elements of a vector.

Some examples to help illustrate this are given below:

```
float4    vf;

float2    low = vf.lo;    // returns vf.xy
float2    high = vf.hi;   // returns vf.zw

float2    even = vf.even; // returns vf.xz
float2    odd = vf.odd;   // returns vf.yw
```

The suffixes `.lo` (or `.even`) and `.hi` (or `.odd`) for a 3-component vector type operate as if the 3-component vector type is a 4-component vector type with the value in the `w` component undefined.

Some examples are given below:

```
float8    vf;
float4    odd = vf.odd;
float4    even = vf.even;
float2    high = vf.even.hi;
float2    low = vf.odd.lo;

// interleave L+R stereo stream
float4    left, right;
float8    interleaved;
interleaved.even = left;
interleaved.odd = right;

// deinterleave
left = interleaved.even;
right = interleaved.odd;

// transpose a 4x4 matrix
void transpose( float4 m[4] )
{
    // read matrix into a float16 vector
    float16 x = (float16)( m[0], m[1], m[2], m[3] );
    float16 t;

    //transpose
    t.even = x.lo;
    t.odd = x.hi;
    x.even = t.lo;
    x.odd = t.hi;
    //write back
    m[0] = x.lo.lo; // { m[0][0], m[1][0], m[2][0], m[3][0] }
    m[1] = x.lo.hi; // { m[0][1], m[1][1], m[2][1], m[3][1] }
```

```

        m[2] = x.hi.lo;    // { m[0][2], m[1][2], m[2][2], m[3][2] }
        m[3] = x.hi.hi;    // { m[0][3], m[1][3], m[2][3], m[3][3] }
    }

float3      vf = (float3)(1.0f, 2.0f, 3.0f);
float2      low = vf.lo; // (1.0f, 2.0f);
float2      high = vf.hi; // (3.0f, undefined);

```

It is an error to take the address of a vector element and will result in a compilation error. For example:

```

float8      vf;

float       *f = &vf.x;           // is illegal
float2      *f2 = &vf.s07;        // is illegal

float4      *odd = &vf.odd;       // is illegal
float4      *even = &vf.even;     // is illegal
float2      *high = &vf.even.hi;  // is illegal
float2      *low = &vf.odd.lo;    // is illegal

```

6.1.8 Aliasing Rules

OpenCL C programs shall comply with the C99 type-based aliasing rules (defined in *section 6.5, item 7* of the C99 specification). The OpenCL C built-in vector data types are considered aggregate²⁶ types for the purpose of applying these aliasing rules.

6.1.9 Keywords

The following names are reserved for use as keywords in OpenCL C and shall not be used otherwise.

- ✚ Names reserved as keywords by C99.
- ✚ OpenCL C data types defined in *tables 6.2, 6.3 and 6.4*.
- ✚ Address space qualifiers: `__global`, `global`, `__local`, `local`, `__constant`, `constant`, `__private` and `private`.
- ✚ Function qualifiers: `__kernel` and `kernel`.
- ✚ Access qualifiers: `__read_only`, `read_only`, `__write_only`, `write_only`, `__read_write` and `read_write`.

²⁶ That is, for the purpose of applying type-based aliasing rules, a built-in vector data type will be considered equivalent to the corresponding array type.

6.2 Conversions and Type Casting

6.2.1 Implicit Conversions

Implicit conversions between scalar built-in types defined in *table 6.1* (except `void` and `half`²⁷) are supported. When an implicit conversion is done, it is not just a re-interpretation of the expression's value but a conversion of that value to an equivalent value in the new type. For example, the integer value 5 will be converted to the floating-point value 5.0.

Implicit conversions between built-in vector data types are disallowed.

Implicit conversions for pointer types follow the rules described in the C99 specification.

6.2.2 Explicit Casts

Standard typecasts for built-in scalar data types defined in *table 6.1* will perform appropriate conversion (except `void` and `half`²⁸). In the example below:

```
float    f = 1.0f;
int      i = (int)f;
```

`f` stores `0x3F800000` and `i` stores `0x1` which is the floating-point value `1.0f` in `f` converted to an integer value.

Explicit casts between vector types are not legal. The examples below will generate a compilation error.

```
int4     i;
uint4    u = (uint4) i; ← not allowed

float4   f;
int4     i = (int4) f; ← not allowed

float4   f;
int8     i = (int8) f; ← not allowed
```

Scalar to vector conversions may be performed by casting the scalar to the desired vector data type. Type casting will also perform appropriate arithmetic conversion. The round to zero rounding mode will be used for conversions to built-in integer vector types. The default rounding mode will be used for conversions to floating-point vector types. When casting a `bool` to a vector integer data type, the vector components will be set to -1 (i.e. all bits set) if the

²⁷ Unless the `cl_khr_fp16` extension is supported.

²⁸ Unless the `cl_khr_fp16` extension is supported.

bool value is *true* and 0 otherwise.

Below are some correct examples of explicit casts.

```
float f = 1.0f;
float4 va = (float4)f;

// va is a float4 vector with elements (f, f, f, f).

uchar u = 0xFF;
float4 vb = (float4)u;

// vb is a float4 vector with elements((float)u, (float)u,
//                                     (float)u, (float)u).

float f = 2.0f;
int2 vc = (int2)f;

// vc is an int2 vector with elements ((int)f, (int)f).

uchar4 vtrue = (uchar4>true;

// vtrue is a uchar4 vector with elements (0xff, 0xff,
//                                         0xff, 0xff).
```

6.2.3 Explicit Conversions

Explicit conversions may be performed using the

```
convert_destType(sourceType)
```

suite of functions. These provide a full set of type conversions between supported types (see *sections 6.1.1, 6.1.2 and 6.1.3*) except for the following types: `bool`, `half`, `size_t`, `ptrdiff_t`, `intptr_t`, `uintptr_t`, and `void`.

The number of elements in the source and destination vectors must match.

In the example below:

```
uchar4 u;
int4 c = convert_int4(u);
```

`convert_int4` converts a `uchar4` vector `u` to an `int4` vector `c`.

```
float f;
int i = convert_int(f);
```

`convert_int` converts a `float` scalar `f` to an `int` scalar `i`.

The behavior of the conversion may be modified by one or two optional modifiers that specify saturation for out-of-range inputs and rounding behavior.

The full form of the scalar `convert` function is:

```
destType convert_destType<_sat><_roundingMode> (sourceType)
```

The full form of the vector `convert` function is:

```
destTypen convert_destTypen<_sat><_roundingMode> (sourceTypen)
```

6.2.3.1 Data Types

Conversions are available for the following scalar types: `char`, `uchar`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `float`, and built-in vector types derived therefrom. The operand and result type must have the same number of elements. The operand and result type may be the same type in which case the conversion has no effect on the type or value of an expression.

Conversions between integer types follow the conversion rules specified in *sections 6.3.1.1 and 6.3.1.3* of the C99 specification except for out-of-range behavior and saturated conversions which are described in *section 6.2.3.3* below.

6.2.3.2 Rounding Modes

Conversions to and from floating-point type shall conform to IEEE-754 rounding rules. Conversions may have an optional rounding mode modifier described in *table 6.6*.

| Modifier | Rounding Mode Description |
|-----------------------|--|
| <code>_rte</code> | Round to nearest even |
| <code>_rtz</code> | Round toward zero |
| <code>_rtp</code> | Round toward positive infinity |
| <code>_rtn</code> | Round toward negative infinity |
| no modifier specified | Use the default rounding mode for this destination type, <code>_rtz</code> for conversion to integers or the default rounding mode for conversion to floating-point types. |

Table 6.6 *Rounding Modes*

By default, conversions to integer type use the `_rtz` (round toward zero) rounding mode and conversions to floating-point type²⁹ use the default rounding mode. The only default floating-

²⁹ For conversions to floating-point format, when a finite source value exceeds the maximum representable finite

point rounding mode supported is round to nearest even i.e the default rounding mode will be `_rte` for floating-point types.

6.2.3.3 Out-of-Range Behavior and Saturated Conversions

When the conversion operand is either greater than the greatest representable destination value or less than the least representable destination value, it is said to be out-of-range. The result of out-of-range conversion is determined by the conversion rules specified by the C99 specification in *section 6.3*. When converting from a floating-point type to integer type, the behavior is implementation-defined.

Conversions to integer type may opt to convert using the optional saturated mode by appending the `_sat` modifier to the conversion function name. When in saturated mode, values that are outside the representable range shall clamp to the nearest representable value in the destination format. (NaN should be converted to 0).

Conversions to floating-point type shall conform to IEEE-754 rounding rules. The `_sat` modifier may not be used for conversions to floating-point formats.

6.2.3.4 Explicit Conversion Examples

Example 1:

```
short4  s;

// negative values clamped to 0
ushort4  u = convert_ushort4_sat( s );

// values > CHAR_MAX converted to CHAR_MAX
// values < CHAR_MIN converted to CHAR_MIN
char4  c = convert_char4_sat( s );
```

Example 2:

```
float4  f;

// values implementation defined for
// f > INT_MAX, f < INT_MIN or NaN
int4  i = convert_int4( f );

// values > INT_MAX clamp to INT_MAX, values < INT_MIN clamp
// to INT_MIN. NaN should produce 0.
```

floating-point destination value, the rounding mode will affect whether the result is the maximum finite floating-point value or infinity of same sign as the source value, per IEEE-754 rules for rounding.

```

// The _rtz rounding mode is used to produce the integer values.
int4 i2 = convert_int4_sat( f );

// similar to convert_int4, except that floating-point values
// are rounded to the nearest integer instead of truncated
int4 i3 = convert_int4_rte( f );

// similar to convert_int4_sat, except that floating-point values
// are rounded to the nearest integer instead of truncated
int4 i4 = convert_int4_sat_rte( f );

```

Example 3:

```

int4 i;

// convert ints to floats using the default rounding mode.
float4 f = convert_float4( i );

// convert ints to floats. integer values that cannot
// be exactly represented as floats should round up to the
// next representable float.
float4 f = convert_float4_rtp( i );

```

6.2.4 Reinterpreting Data As Another Type

It is frequently necessary to reinterpret bits in a data type as another data type in OpenCL. This is typically required when direct access to the bits in a floating-point type is needed, for example to mask off the sign bit or make use of the result of a vector relational operator (see *section 6.3.d*) on floating-point data³⁰. Several methods to achieve this (non-) conversion are frequently practiced in C, including pointer aliasing, unions and memcpy. Of these, only memcpy is strictly correct in C99. Since OpenCL does not provide **memcpy**, other methods are needed.

6.2.4.1 Reinterpreting Types Using Unions

The OpenCL language extends the union to allow the program to access a member of a union object using a member of a different type. The relevant bytes of the representation of the object are treated as an object of the type used for the access. If the type used for access is larger than the representation of the object, then the value of the additional bytes is undefined.

³⁰ In addition, some other extensions to the C language designed to support particular vector ISA (e.g. AltiVec™, CELL Broadband Engine™ Architecture) use such conversions in conjunction with swizzle operators to achieve type unconversion. So as to support legacy code of this type, `as_typen()` allows conversions between vectors of the same size but different numbers of elements, even though the behavior of this sort of conversion is not likely to be portable except to other OpenCL implementations for the same hardware architecture. AltiVec™ is a trademark of Motorola Inc. Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc.

Examples:

```
union{ float f; uint u; double d31;} u;

u.u = 1;           // u.f contains 2**-149.  u.d is undefined --
                  // depending on endianness the low or high half
                  // of d is unknown

u.f = 1.0f;       // u.u contains 0x3f800000, u.d contains an
                  // undefined value -- depending on endianness
                  // the low or high half of d is unknown

u.d = 1.0;        // u.u contains 0x3ff00000 (big endian) or 0
                  // (little endian). u.f contains either 0x1.ep0f
                  // (big endian) or 0.0f (little endian)
```

6.2.4.2 Reinterpreting Types Using `as_type()` and `as_type_n()`

All data types described in tables 6.1 and 6.2 (except `bool`, `half`³² and `void`) may be also reinterpreted as another data type of the same size using the `as_type()` operator for scalar data types and the `as_type_n()` operator³³ for vector data types. When the operand and result type contain the same number of elements, the bits in the operand shall be returned directly without modification as the new type. The usual type promotion for function arguments shall not be performed.

For example, `as_float(0x3f800000)` returns `1.0f`, which is the value that the bit pattern `0x3f800000` has if viewed as an IEEE-754 single precision value.

When the operand and result type contain a different number of elements, the result shall be implementation-defined except if the operand is a 4-component vector and the result is a 3-component vector. In this case, the bits in the operand shall be returned directly without modification as the new type. That is, a conforming implementation shall explicitly define a behavior, but two conforming implementations need not have the same behavior when the number of elements in the result and operand types does not match. The implementation may

³¹ Only if double precision is supported.

³² Unless the `cl_khr_fp16` extension is supported.

³³ While the union is intended to reflect the organization of data in memory, the `as_type()` and `as_type_n()` constructs are intended to reflect the organization of data in register. The `as_type()` and `as_type_n()` constructs are intended to compile to no instructions on devices that use a shared register file designed to operate on both the operand and result types. Note that while differences in memory organization are expected to largely be limited to those arising from endianness, the register based representation may also differ due to size of the element in register. (For example, an architecture may load a char into a 32-bit register, or a char vector into a SIMD vector register with fixed 32-bit element size.) If the element count does not match, then the implementation should pick a data representation that most closely matches what would happen if an appropriate result type operator was applied to a register containing data of the source type. If the number of elements matches, then the `as_type_n()` should faithfully reproduce the behavior expected from a similar data type reinterpretation using memory/unions. So, for example if an implementation stores all single precision data as double in register, it should implement `as_int(float)` by first downconverting the double to single precision and then (if necessary) moving the single precision bits to a register suitable for operating on integer data. If data stored in different address spaces do not have the same endianness, then the "dominant endianness" of the device should prevail.

define the result to contain all, some or none of the original bits in whatever order it chooses. It is an error to use `as_type()` or `as_typen()` operator to reinterpret data to a type of a different number of bytes.

Examples:

```
float f = 1.0f;
uint u = as_uint(f); // Legal. Contains: 0x3f800000

float4 f = (float4)(1.0f, 2.0f, 3.0f, 4.0f);
// Legal. Contains:
// (int4)(0x3f800000, 0x40000000, 0x40400000, 0x40800000)
int4 i = as_int4(f);

float4 f, g;
int4 is_less = f < g;

// Legal. f[i] = f[i] < g[i] ? f[i] : 0.0f
f = as_float4(as_int4(f) & is_less);

int i;
// Legal. Result is implementation-defined.
short2 j = as_short2(i);

int4 i;
// Legal. Result is implementation-defined.
short8 j = as_short8(i);

float4 f;
// Error. Result and operand have different sizes
double4 g = as_double434(f);

float4 f;
// Legal. g.xyz will have same values as f.xyz. g.w is undefined
float3 g = as_float3(f);
```

6.2.5 Pointer Casting

Pointers to old and new types may be cast back and forth to each other. Casting a pointer to a new type represents an unchecked assertion that the address is correctly aligned. The developer will also need to know the endianness of the OpenCL device and the endianness of the data to determine how the scalar and vector data elements are stored in memory.

³⁴ Only if double precision is supported.

6.2.6 Usual Arithmetic Conversions

Many operators that expect operands of arithmetic type cause conversions and yield result types in a similar way. The purpose is to determine a common real type for the operands and result. For the specified operands, each operand is converted, without change of type domain, to a type whose corresponding real type is the common real type. For this purpose, all vector types shall be considered to have higher conversion ranks than scalars. Unless explicitly stated otherwise, the common real type is also the corresponding real type of the result, whose type domain is the type domain of the operands if they are the same, and complex otherwise. This pattern is called the usual arithmetic conversions. If the operands are of more than one vector type, then an error shall occur. Implicit conversions between vector types are not permitted, per *section 6.2.1*.

Otherwise, if there is only a single vector type, and all other operands are scalar types, the scalar types are converted to the type of the vector element, then widened into a new vector containing the same number of elements as the vector, by duplication of the scalar value across the width of the new vector. An error shall occur if any scalar operand has greater rank than the type of the vector element. For this purpose, the rank order defined as follows:

1. The rank of a floating-point type is greater than the rank of another floating-point type, if the first floating-point type can exactly represent all numeric values in the second floating-point type. (For this purpose, the encoding of the floating-point value is used, rather than the subset of the encoding usable by the device.)
2. The rank of any floating-point type is greater than the rank of any integer type.
3. The rank of an integer type is greater than the rank of an integer type with less precision.
4. The rank of an unsigned integer type is **greater than** the rank of a signed integer type with the same precision.³⁵
5. The rank of the bool type is less than the rank of any other type.
6. The rank of an enumerated type shall equal the rank of the compatible integer type.
7. For all types, T1, T2 and T3, if T1 has greater rank than T2, and T2 has greater rank than T3, then T1 has greater rank than T3.

Otherwise, if all operands are scalar, the usual arithmetic conversions apply, per *section 6.3.1.8* of the C99 standard.

NOTE: Both the standard orderings in *sections 6.3.1.8* and *6.3.1.1* of C99 were examined and rejected. Had we used integer conversion rank here, `int4 + 0U` would have been legal and had `int4` return type. Had we used standard C99 usual arithmetic conversion rules for scalars, then the standard integer promotion would have been performed on vector integer element types and `short8 + char` would either have return type of `int8` or be illegal.

³⁵ This is different from the standard integer conversion rank described in C99 TC2, *section 6.3.1.1*.

6.3 Operators

- a. The arithmetic operators add (+), subtract (-), multiply (*) and divide (/) operate on built-in integer and floating-point scalar, and vector data types. The remainder (%) operates on built-in integer scalar and integer vector data types. All arithmetic operators return result of the same built-in type (integer or floating-point) as the type of the operands, after operand type conversion. After conversion, the following cases are valid:

- ✚ The two operands are scalars. In this case, the operation is applied, resulting in a scalar.
- ✚ One operand is a scalar, and the other is a vector. In this case, the scalar may be subject to the usual arithmetic conversion to the element type used by the vector operand. The scalar type is then widened to a vector that has the same number of components as the vector operand. The operation is done component-wise resulting in the same size vector.
- ✚ The two operands are vectors of the same type. In this case, the operation is done component-wise resulting in the same size vector.

All other cases of implicit conversions are illegal. Division on integer types which results in a value that lies outside of the range bounded by the maximum and minimum representable values of the integer type will not cause an exception but will result in an unspecified value. A divide by zero with integer types does not cause an exception but will result in an unspecified value. Division by zero for floating-point types will result in \pm infinity or NaN as prescribed by the IEEE-754 standard. Use the built-in functions **dot** and **cross** to get, respectively, the vector dot product and the vector cross product.

- b. The arithmetic unary operators (+ and -) operate on built-in scalar and vector types.
- c. The arithmetic post- and pre-increment and decrement operators (-- and ++) operate on built-in scalar and vector types except the built-in scalar and vector float types³⁶. All unary operators work component-wise on their operands. These result with the same type they operated on. For post- and pre-increment and decrement, the expression must be one that could be assigned to (an l-value). Pre-increment and pre-decrement add or subtract 1 to the contents of the expression they operate on, and the value of the pre-increment or pre-decrement expression is the resulting value of that modification. Post-increment and post-decrement expressions add or subtract 1 to the contents of the expression they operate on, but the resulting expression has the expression's value before the post-increment or post-decrement was executed.

³⁶ The pre- and post- increment operators may have unexpected behavior on floating-point values and are therefore not supported for floating-point scalar and vector built-in types. For example, if variable *a* has type float and holds the value 0x1.0p25f, then *a*++ returns 0x1.0p25f. Also, (*a*++)-- is not guaranteed to return *a*, if *a* has fractional value. In non-default rounding modes, (*a*++)-- may produce the same result as *a*++ or *a*-- for large *a*.

d. The relational operators³⁷ greater than ($>$), less than ($<$), greater than or equal ($>=$), and less than or equal ($<=$) operate on scalar and vector types. All relational operators result in an integer type. After operand type conversion, the following cases are valid:

- ✚ The two operands are scalars. In this case, the operation is applied, resulting in an `int` scalar.
- ✚ One operand is a scalar, and the other is a vector. In this case, the scalar may be subject to the usual arithmetic conversion to the element type used by the vector operand. The scalar type is then widened to a vector that has the same number of components as the vector operand. The operation is done component-wise resulting in the same size vector.
- ✚ The two operands are vectors of the same type. In this case, the operation is done component-wise resulting in the same size vector.

All other cases of implicit conversions are illegal.

The result is a scalar signed integer of type `int` if the source operands are scalar and a vector signed integer type of the same size as the source operands if the source operands are vector types. Vector source operands of type `char n` and `uchar n` return a `char n` result; vector source operands of type `short n` and `ushort n` return a `short n` result; vector source operands of type `int n` , `uint n` and `float n` return an `int n` result; vector source operands of type `long n` , `ulong n` and `double n` return a `long n` result. For scalar types, the relational operators shall return 0 if the specified relation is *false* and 1 if the specified relation is *true*. For vector types, the relational operators shall return 0 if the specified relation is *false* and -1 (i.e. all bits set) if the specified relation is *true*. The relational operators always return 0 if either argument is not a number (NaN).

e. The equality operators³⁸ equal ($==$), and not equal ($!=$) operate on built-in scalar and vector types. All equality operators result in an integer type. After operand type conversion, the following cases are valid:

- ✚ The two operands are scalars. In this case, the operation is applied, resulting in a scalar.
- ✚ One operand is a scalar, and the other is a vector. In this case, the scalar may be subject to the usual arithmetic conversion to the element type used by the vector operand. The scalar type is then widened to a vector that has the same number of components as the vector operand. The operation is done component-wise resulting

³⁷ To test whether any or all elements in the result of a vector relational operator test true, for example to use in the context in an `if ()` statement, please see the **any** and **all** builtins in *section 6.11.6*.

³⁸ To test whether any or all elements in the result of a vector equality operator test true, for example to use in the context in an `if ()` statement, please see the **any** and **all** builtins in *section 6.11.6*.

in the same size vector.

- ✚ The two operands are vectors of the same type. In this case, the operation is done component-wise resulting in the same size vector.

All other cases of implicit conversions are illegal.

The result is a scalar signed integer of type `int` if the source operands are scalar and a vector signed integer type of the same size as the source operands if the source operands are vector types. Vector source operands of type `char n` and `uchar n` return a `char n` result; vector source operands of type `short n` and `ushort n` return a `short n` result; vector source operands of type `int n` , `uint n` and `float n` return an `int n` result; vector source operands of type `long n` , `ulong n` and `double n` return a `long n` result.

For scalar types, the equality operators return 0 if the specified relation is *false* and return 1 if the specified relation is *true*. For vector types, the equality operators shall return 0 if the specified relation is *false* and -1 (i.e. all bits set) if the specified relation is *true*. The equality operator `equal (==)` returns 0 if one or both arguments are not a number (NaN). The equality operator `not equal (!=)` returns 1 (for scalar source operands) or -1 (for vector source operands) if one or both arguments are not a number (NaN).

- The bitwise operators `and (&)`, `or (|)`, `exclusive or (^)`, `not (~)` operate on all scalar and vector built-in types except the built-in scalar and vector float types. For vector built-in types, the operators are applied component-wise. If one operand is a scalar and the other is a vector, the scalar may be subject to the usual arithmetic conversion to the element type used by the vector operand. The scalar type is then widened to a vector that has the same number of components as the vector operand. The operation is done component-wise resulting in the same size vector.
- The logical operators `and (&&)`, `or (||)` operate on all scalar and vector built-in types. For scalar built-in types only, `and (&&)` will only evaluate the right hand operand if the left hand operand compares unequal to 0. For scalar built-in types only, `or (||)` will only evaluate the right hand operand if the left hand operand compares equal to 0. For built-in vector types, both operands are evaluated and the operators are applied component-wise. If one operand is a scalar and the other is a vector, the scalar may be subject to the usual arithmetic conversion to the element type used by the vector operand. The scalar type is then widened to a vector that has the same number of components as the vector operand. The operation is done component-wise resulting in the same size vector.

The logical operator `exclusive or (^)` is reserved.

The result is a scalar signed integer of type `int` if the source operands are scalar and a vector signed integer type of the same size as the source operands if the source operands are vector types. Vector source operands of type `char n` and `uchar n` return a `char n` result; vector source operands of type `short n` and `ushort n` return a `short n` result; vector source

operands of type `intn`, `uintn` and `floatn` return an `intn` result; vector source operands of type `longn`, `ulongn` and `doublen` return a `longn` result.

For scalar types, the logical operators shall return 0 if the result of the operation is *false* and 1 if the result is *true*. For vector types, the logical operators shall return 0 if the result of the operation is *false* and -1 (i.e. all bits set) if the result is *true*.

- h. The logical unary operator not (!) operates on all scalar and vector built-in types. For built-in vector types, the operators are applied component-wise.

The result is a scalar signed integer of type `int` if the source operands are scalar and a vector signed integer type of the same size as the source operands if the source operands are vector types. Vector source operands of type `charn` and `ucharn` return a `charn` result; vector source operands of type `shortn` and `ushortn` return a `shortn` result; vector source operands of type `intn`, `uintn` and `floatn` return an `intn` result; vector source operands of type `longn`, `ulongn` and `doublen` return a `longn` result.

For scalar types, the result of the logical unary operator is 0 if the value of its operand compares unequal to 0, and 1 if the value of its operand compares equal to 0. For vector types, the unary operator shall return a 0 if the value of its operand compares unequal to 0, and -1 (i.e. all bits set) if the value of its operand compares equal to 0.

- i. The ternary selection operator (`? :`) operates on three expressions (`exp1 ? exp2 : exp3`). This operator evaluates the first expression `exp1`, which can be a scalar or vector result except float. If the result is a scalar value then it selects to evaluate the second expression if the result compares unequal to 0, otherwise it selects to evaluate the third expression. If the result is a vector value, then this is equivalent to calling `select(exp3, exp2, exp1)`. The `select` function is described in *table 6.14*. The second and third expressions can be any type, as long their types match, or there is a conversion in *section 6.2.1 Implicit Conversions* that can be applied to one of the expressions to make their types match, or one is a vector and the other is a scalar and the scalar may be subject to the usual arithmetic conversion to the element type used by the vector operand and widened to the same type as the vector type. This resulting matching type is the type of the entire expression.
- j. The operators right-shift (`>>`), left-shift (`<<`) operate on all scalar and vector built-in types except the built-in scalar and vector float types. For built-in vector types, the operators are applied component-wise. For the right-shift (`>>`), left-shift (`<<`) operators, the rightmost operand must be a scalar if the first operand is a scalar, and the rightmost operand can be a vector or scalar if the first operand is a vector.

The result of `E1 << E2` is `E1` left-shifted by $\log_2(N)$ least significant bits in `E2` viewed as an unsigned integer value, where `N` is the number of bits used to represent the data type of `E1` after integer promotion³⁹, if `E1` is a scalar, or the number of bits used to represent the

³⁹ Integer promotion is described in ISO/IEC 9899:1999 in *section 6.3.1.1*.

type of E1 elements, if E1 is a vector. The vacated bits are filled with zeros.

The result of `E1 >> E2` is E1 right-shifted by $\log_2(N)$ least significant bits in E2 viewed as an unsigned integer value, where N is the number of bits used to represent the data type of E1 after integer promotion, if E1 is a scalar, or the number of bits used to represent the type of E1 elements, if E1 is a vector. If E1 has an unsigned type or if E1 has a signed type and a nonnegative value, the vacated bits are filled with zeros. If E1 has a signed type and a negative value, the vacated bits are filled with ones.

- k. The `sizeof` operator yields the size (in bytes) of its operand, including any padding bytes (refer to *section 6.1.5*) needed for alignment, which may be an expression or the parenthesized name of a type. The size is determined from the type of the operand. The result is of type `size_t`. If the type of the operand is a variable length array⁴⁰ type, the operand is evaluated; otherwise, the operand is not evaluated and the result is an integer constant.

When applied to an operand that has type `char`, `uchar`, the result is 1. When applied to an operand that has type `short`, `ushort`, or `half` the result is 2. When applied to an operand that has type `int`, `uint` or `float`, the result is 4. When applied to an operand that has type `long`, `ulong` or `double`, the result is 8. When applied to an operand that is a vector type, the result⁴¹ is number of components * size of each scalar component. When applied to an operand that has array type, the result is the total number of bytes in the array. When applied to an operand that has structure or union type, the result is the total number of bytes in such an object, including internal and trailing padding. The `sizeof` operator shall not be applied to an expression that has function type or an incomplete type, to the parenthesized name of such a type, or to an expression that designates a bit-field struct member⁴².

The behavior of applying the `sizeof` operator to the `bool`, `image2d_t`, `image3d_t`, `image2d_array_t`, `image1d_t`, `image1d_buffer_t` or `image1d_array_t`, `sampler_t` and `event_t` types is implementation-defined.

- l. The comma (,) operator operates on expressions by returning the type and value of the right-most expression in a comma separated list of expressions. All expressions are evaluated, in order, from left to right.
- m. The unary (*) operator denotes indirection. If the operand points to an object, the result is an lvalue designating the object. If the operand has type “pointer to *type*”, the result has type “*type*”. If an invalid value has been assigned to the pointer, the behavior of the unary *

⁴⁰ Variable length arrays are not supported in OpenCL 1.1. Refer to *section 6.9.d*.

⁴¹ Except for 3-component vectors whose size is defined as 4 * size of each scalar component.

⁴² Bit-field struct members are not supported in OpenCL 1.1. Refer to *section 6.9.c*.

operator is undefined⁴³.

- n. The unary (&) operator returns the address of its operand. If the operand has type “*type*”, the result has type “pointer to *type*”. If the operand is the result of a unary * operator, neither that operator nor the & operator is evaluated and the result is as if both were omitted, except that the constraints on the operators still apply and the result is not an lvalue. Similarly, if the operand is the result of a [] operator, neither the & operator nor the unary * that is implied by the [] is evaluated and the result is as if the & operator were removed and the [] operator were changed to a + operator. Otherwise, the result is a pointer to the object designated by its operand⁴⁴.
- o. Assignments of values to variable names are done with the assignment operator (=), like

```
lvalue = expression
```

The assignment operator stores the value of *expression* into *lvalue*. The *expression* and *lvalue* must have the same type, or the expression must have a type in *table 6.1*, in which case an implicit conversion will be done on the expression before the assignment is done.

If *expression* is a scalar type and *lvalue* is a vector type, the scalar is converted to the element type used by the vector operand. The scalar type is then widened to a vector that has the same number of components as the vector operand. The operation is done component-wise resulting in the same size vector.

Any other desired type-conversions must be specified explicitly. L-values must be writable. Variables that are built-in types, entire structures or arrays, structure fields, l-values with the field selector (.) applied to select components or swizzles without repeated fields, l-values within parentheses, and l-values dereferenced with the array subscript operator ([]) are all l-values. Other binary or unary expressions, function names, swizzles with repeated fields, and constants cannot be l-values. The ternary operator (?:) is also not allowed as an l-value.

The order of evaluation of the operands is unspecified. If an attempt is made to modify the result of an assignment operator or to access it after the next sequence point, the behavior is undefined. Other assignment operators are the assignments add into (+=), subtract from (-=), multiply into (*=), divide into (/=), modulus into (%=), left shift by (<<=), right shift by (>>=), and into (&=), inclusive or into (|=), and exclusive or into (^=).

The expression

```
lvalue op= expression
```

⁴³ Among the invalid values for dereferencing a pointer by the unary * operator are a null pointer, an address inappropriately aligned for the type of object pointed to, and the address of an object after the end of its lifetime. If ***P** is an lvalue and **T** is the name of an object pointer type, *(**T**)**P** is an lvalue that has a type compatible with that to which **T** points.

⁴⁴ Thus, &***E** is equivalent to **E** (even if **E** is a null pointer), and &(**E1**[**E2**]) to ((**E1**)+(**E2**)). It is always true that if **E** is an lvalue that is a valid operand of the unary & operator, *&**E** is an lvalue equal to **E**.

is equivalent to

```
lvalue = lvalue op expression
```

and the l-value and expression must satisfy the requirements for both operator *op* and assignment (=).

Note: Except for the `sizeof` operator, the `half` data type cannot be used with any of the operators described in this section.

6.4 Vector Operations

Vector operations are component-wise. Usually, when an operator operates on a vector, it is operating independently on each component of the vector, in a component-wise fashion.

For example,

```
float4    v, u;  
float     f;  
  
v = u + f;
```

will be equivalent to

```
v.x = u.x + f;  
v.y = u.y + f;  
v.z = u.z + f;  
v.w = u.w + f;
```

And

```
float4    v, u, w;  
  
w = v + u;
```

will be equivalent to

```
w.x = v.x + u.x;  
w.y = v.y + u.y;  
w.z = v.z + u.z;  
w.w = v.w + u.w;
```

and likewise for most operators and all integer and floating-point vector types.

6.5 Address Space Qualifiers

OpenCL implements the following disjoint address spaces: `__global`, `__local`, `__constant` and `__private`. The address space qualifier may be used in variable declarations to specify the region of memory that is used to allocate the object. The C syntax for type qualifiers is extended in OpenCL to include an address space name as a valid type qualifier. If the type of an object is qualified by an address space name, the object is allocated in the specified address name; otherwise, the object is allocated in the generic address space.

The address space names without the `__` prefix i.e. `global`, `local`, `constant` and `private` may be substituted for the corresponding address space names with the `__` prefix.

The generic address space name for arguments to a function in a program, or local variables of a function is `__private`. All function arguments shall be in the `__private` address space.

`__kernel` function arguments declared to be a pointer or an array of a type can point to one of the following address spaces only: `__global`, `__local` or `__constant`. A pointer to address space A can only be assigned to a pointer to the same address space A. Casting a pointer to address space A to a pointer to address space B is illegal.

There is no generic address space name for program scope variables. All program scope variables must be declared in the `__constant` address space.

Examples:

```
// declares a pointer p in the __private address space that
// points to an int object in address space __global
__global int *p;

// declares an array of 4 floats in the __private address space.
float x[4];
```

There is no address space for function return values. Using an address space qualifier in a function return type declaration will generate a compilation error, unless the return type is declared as a pointer type and the qualifier is used on the points-to address space.

Examples:

```
__private int f() { ... } // should generate an error
__local int *f() { ... } // allowed
__local int * __private f() { ... }; // should generate an error.
```


6.5.1 `__global` (or `global`)

The `__global` or `global` address space name is used to refer to memory objects (buffer or image objects) allocated from the global memory pool.

A buffer memory object can be declared as a pointer to a scalar, vector or user-defined struct. This allows the kernel to read and/or write any location in the buffer.

The actual size of the array memory object is determined when the memory object is allocated via appropriate API calls in the host code.

Some examples are:

```
__global float4  *color;    // An array of float4 elements
typedef struct {
    float a[3];
    int   b[2];
} foo_t;
__global foo_t   *my_info; // An array of foo_t elements.
```

As image objects are always allocated from the global address space, the `__global` or `global` qualifier should not be specified for image types. The elements of an image object cannot be directly accessed. Built-in functions to read from and write to an image object are provided.

The `const` qualifier can also be used with the `__global` qualifier to specify a read-only buffer memory object.

6.5.2 `__local` (or `local`)

The `__local` or `local` address space name is used to describe variables that need to be allocated in local memory and are shared by all work-items of a work-group. Pointers to the `__local` address space are allowed as arguments to functions (including kernel functions). Variables declared in the `__local` address space inside a kernel function must occur at kernel function scope.

Some examples of variables allocated in the `__local` address space inside a kernel function are:

```
__kernel void my_func(...)
{
    __local float   a;    // A single float allocated
                        // in local address space

    __local float   b[10]; // An array of 10 floats
```

```

// allocated in local address space.

if (...)
{
    // example of variable in __local address space but not
    // declared at __kernel function scope.
    __local float    c;    ← not allowed.
}
}

```

Variables allocated in the **__local** address space inside a kernel function cannot be initialized.

```

__kernel void my_func(...)
{
    __local float    a = 1;    ← not allowed

    __local float    b;
    b = 1;                  ← allowed
}

```

NOTE: Variables allocated in the **__local** address space inside a kernel function are allocated for each work-group executing the kernel and exist only for the lifetime of the work-group executing the kernel.

6.5.3 **__constant** (or constant)

The **__constant** or **constant** address space name is used to describe variables allocated in global memory and which are accessed inside a kernel(s) as read-only variables. These read-only variables can be accessed by all (global) work-items of the kernel during its execution. Pointers to the **__constant** address space are allowed as arguments to functions (including kernel functions) and for variables declared inside functions.

All string literal storage shall be in the **__constant** address space.

NOTE: Each argument to a kernel that is a pointer to the **__constant** address space is counted separately towards the maximum number of such arguments, defined as `CL_DEVICE_MAX_CONSTANT_ARGS` in *table 4.3*.

Variables in the program scope must be declared in the **__constant** address space. Variables in the outermost scope of kernel functions can be declared in the **__constant** address space. These variables are required to be initialized and the values used to initialize these variables must be a compile time constant. Writing to such a variable results in a compile-time error.

Implementations are not required to aggregate these declarations into the fewest number of constant arguments. This behavior is implementation defined.

Thus portable code must conservatively assume that each variable declared inside a function or in program scope allocated in the `__constant` address space counts as a separate constant argument.

6.5.4 `__private` (or `private`)

Variables inside a kernel function not declared with an address space qualifier, all variables inside non-kernel functions, and all function arguments are in the `__private` or `private` address space. Variables declared as pointers are considered to point to the `__private` address space if an address space qualifier is not specified.

The `__global`, `__constant`, `__local`, `__private`, `global`, `constant`, `local` and `private` names are reserved for use as address space qualifiers and shall not be used otherwise.

6.6 Access Qualifiers

Image objects specified as arguments to a kernel can be declared to be read-only or write-only. A kernel cannot read from and write to the same image object. The **`__read_only`** (or **`read_only`**) and **`__write_only`** (or **`write_only`**) qualifiers must be used with image object arguments to declare if the image object is being read or written by a kernel. The default qualifier is **`__read_only`**.

In the following example

```
__kernel void
foo (read_only image2d_t imageA,
     write_only image2d_t imageB)
{
    ...
}
```

`imageA` is a read-only 2D image object, and `imageB` is a write-only 2D image object.

The `__read_only`, `__write_only`, `__read_write`, `read_only`, `write_only` and `read_write` names are reserved for use as access qualifiers and shall not be used otherwise.

6.7 Function Qualifiers

6.7.1 `__kernel` (or `kernel`)

The `__kernel` (or `kernel`) qualifier declares a function to be a kernel that can be executed by an application on an OpenCL device(s). The following rules apply to functions that are declared with this qualifier:

- ✚ It can be executed on the device only
- ✚ It can be called by the host
- ✚ It is just a regular function call if a `__kernel` function is called by another kernel function.

NOTE:

Kernel functions with variables declared inside the function with the `__local` or `local` qualifier can be called by the host using appropriate APIs such as `clEnqueueNDRangeKernel`, and `clEnqueueTask`.

The `__kernel` and `kernel` names are reserved for use as functions qualifiers and shall not be used otherwise.

6.7.2 Optional Attribute Qualifiers

The `__kernel` qualifier can be used with the keyword `__attribute__` to declare additional information about the kernel function as described below.

The optional `__attribute__((vec_type_hint(<type>)))`⁴⁵ is a hint to the compiler and is intended to be a representation of the computational *width* of the `__kernel`, and should serve as the basis for calculating processor bandwidth utilization when the compiler is looking to autovectorize the code. In the `__attribute__((vec_type_hint(<type>)))` qualifier `<type>` is one of the built-in vector types listed in *table 6.2* or the constituent scalar element types. If `vec_type_hint (<type>)` is not specified, the kernel is assumed to have the `__attribute__((vec_type_hint(int)))` qualifier.

For example, where the developer specified a width of `float4`, the compiler should assume that the computation usually uses up to 4 lanes of a float vector, and would decide to merge

⁴⁵ Implicit in autovectorization is the assumption that any libraries called from the `__kernel` must be recompilable at run time to handle cases where the compiler decides to merge or separate workitems. This probably means that such libraries can never be hard coded binaries or that hard coded binaries must be accompanied either by source or some retargetable intermediate representation. This may be a code security question for some.

work-items or possibly even separate one work-item into many threads to better match the hardware capabilities. A conforming implementation is not required to autovectorize code, but shall support the hint. A compiler may autovectorize, even if no hint is provided. If an implementation merges N work-items into one thread, it is responsible for correctly handling cases where the number of global or local work-items in any dimension modulo N is not zero.

Examples:

```
// autovectorize assuming float4 as the
// basic computation width
__kernel __attribute__((vec_type_hint(float4)))
void foo( __global float4 *p ) { .... }

// autovectorize assuming double as the
// basic computation width
__kernel __attribute__((vec_type_hint(double)))
void foo( __global float4 *p ) { .... }

// autovectorize assuming int (default)
// as the basic computation width
__kernel
void foo( __global float4 *p ) { .... }
```

If for example, a `__kernel` function is declared with `__attribute__((vec_type_hint(float4)))` (meaning that most operations in the `__kernel` function are explicitly vectorized using `float4`) and the kernel is running using Intel® Advanced Vector Instructions (Intel® AVX) which implements a 8-float-wide vector unit, the autovectorizer might choose to merge two work-items to one thread, running a second work-item in the high half of the 256-bit AVX register.

As another example, a Power4 machine has two scalar double precision floating-point units with an 6-cycle deep pipe. An autovectorizer for the Power4 machine might choose to interleave six kernels declared with the `__attribute__((vec_type_hint(double2)))` qualifier into one hardware thread, to ensure that there is always 12-way parallelism available to saturate the FPUs. It might also choose to merge 4 or 8 work-items (or some other number) if it concludes that these are better choices, due to resource utilization concerns or some preference for divisibility by 2.

The optional `__attribute__((work_group_size_hint(X, Y, Z)))` is a hint to the compiler and is intended to specify the work-group size that may be used i.e. value most likely to be specified by the *local_work_size* argument to `clEnqueueNDRangeKernel`. For example the `__attribute__((work_group_size_hint(1, 1, 1)))` is a hint to the compiler that the kernel will most likely be executed with a work-group size of 1.

The optional `__attribute__((reqd_work_group_size(X, Y, Z)))` is the work-group size that must be used as the *local_work_size* argument to **clEnqueueNDRangeKernel**. This allows the compiler to optimize the generated code appropriately for this kernel. The optional `__attribute__((reqd_work_group_size(X, Y, Z)))`, if specified, must be (1, 1, 1) if the kernel is executed via **clEnqueueTask**.

If Z is one, the *work_dim* argument to **clEnqueueNDRangeKernel** can be 2 or 3. If Y and Z are one, the *work_dim* argument to **clEnqueueNDRangeKernel** can be 1, 2 or 3.

6.8 Storage-Class Specifiers

The `typedef`, `extern` and `static` storage-class specifiers are supported. The `auto` and `register` storage-class specifiers are not supported.

The `extern` storage-class specifier can only be used for functions (kernel and non-kernel functions) and global variables declared in program scope or variables declared inside a function (kernel and non-kernel functions). The `static` storage-class specifier can only be used for non-kernel functions and global variables declared in program scope.

Examples:

```
extern constant float4 noise_table[256];
static constant float4 color_table[256];

extern kernel void my_foo(image2d_t img);
extern void my_bar(global float *a);

kernel void my_func(image2d_t img, global float *a)
{
    extern constant float4 a;
    static constant float4 b;        // error.
    static float c;                 // error.

    ...
    my_foo(img);
    ...
    my_bar(a);
    ...
}
```


6.9 Restrictions⁴⁶

- a. The use of pointers is somewhat restricted. The following rules apply:
 - ✚ Arguments to kernel functions declared in a program that are pointers must be declared with the `__global`, `__constant` or `__local` qualifier.
 - ✚ A pointer declared with the `__constant`, `__local` or `__global` qualifier can only be assigned to a pointer declared with the `__constant`, `__local` or `__global` qualifier respectively.
 - ✚ Pointers to functions are not allowed.
 - ✚ Arguments to kernel functions in a program cannot be declared as a pointer to a pointer(s). Variables inside a function or arguments to non-kernel functions in a program can be declared as a pointer to a pointer(s).
- b. An image type (`image2d_t`, `image3d_t`, `image2d_array_t`, `image1d_t`, `image1d_buffer_t` or `image1d_array_t`) can only be used as the type of a function argument. An image function argument cannot be modified. Elements of an image can only be accessed using built-in functions described in *section 6.12.14*.

An image type cannot be used to declare a variable, a structure or union field, an array of images, a pointer to an image, or the return type of a function. An image type cannot be used with the `__global`, `__private`, `__local` and `__constant` address space qualifiers. The `image3d_t` type cannot be used with the `__write_only` access qualifier unless the `cl_khr_3d_image_writes` extension is enabled. An image type cannot be used with the `__read_write` access qualifer which is reserved for future use.

The sampler type (`sampler_t`) can only be used as the type of a function argument or a variable declared in the program scope or the outermost scope of a kernel function. The behavior of a sampler variable declared in a non-outermost scope of a kernel function is implementation-defined. A sampler argument or variable cannot be modified.

The sampler type cannot be used to declare a structure or union field, an array of samplers, a pointer to a sampler, or the return type of a function. The sampler type cannot be used with the `__local` and `__global` address space qualifiers.
- c. Bit-field struct members are currently not supported.
- d. Variable length arrays and structures with flexible (or unsized) arrays are not supported.

⁴⁶ Items struckthrough are restrictions in OpenCL 1.0 that are removed in OpenCL 1.1.

- e. Variadic macros and functions are not supported.
- f. The library functions defined in the C99 standard headers `assert.h`, `ctype.h`, `complex.h`, `errno.h`, `fenv.h`, `float.h`, `inttypes.h`, `limits.h`, `locale.h`, `setjmp.h`, `signal.h`, `stdarg.h`, `stdio.h`, `stdlib.h`, `string.h`, `tgmath.h`, `time.h`, `wchar.h` and `wctype.h` are not available and cannot be included by a program.
- g. The `auto` and `register` storage-class specifiers are not supported.
- ~~h. Predefined identifiers are not supported.~~
- i. Recursion is not supported.
- j. The return type of a kernel function must be `void`.
- k. Arguments to kernel functions in a program cannot be declared with the built-in scalar types `bool`, `half`, `size_t`, `ptrdiff_t`, `intptr_t`, and `uintptr_t` or a struct and/or union that contain fields declared to be one of these built-in scalar types. The size in bytes of these types except `half` are implementation-defined and in addition can also be different for the OpenCL device and the host processor making it difficult to allocate buffer objects to be passed as arguments to a kernel declared as pointer to these types. `half` is not supported as `half` can be used as a storage format⁴⁷ only and is not a data type on which floating-point arithmetic can be performed.
- l. Whether or not irreducible control flow is illegal is implementation defined.
- ~~m. Built-in types that are less than 32-bits in size i.e. `char`, `uchar`, `char2`, `uchar2`, `short`, `ushort`, and `half` have the following restriction:~~

~~Writes to a pointer (or arrays) of type `char`, `uchar`, `char2`, `uchar2`, `short`, `ushort`, and `half` or to elements of a struct that are of type `char`, `uchar`, `char2`, `uchar2`, `short` and `ushort` are not supported. Refer to [section 9.9](#) for additional information.~~

~~The kernel example below shows what memory operations are not supported on built-in types less than 32-bits in size.~~

```
kernel void
do_proc (__global char *pA, short b,
         __global short *pB)
{
    char x[100];
```

⁴⁷ Unless the `cl_khr_fp16` extension is supported.

```

__private char *px = x;
int id = (int)get_global_id(0);
short f;

f = pB[id] + b; ← is allowed

px[1] = pA[1]; ← error. px cannot be written.

pB[id] = b; ← error. pB cannot be written
}

```

- n. Arguments to kernel functions in a program cannot be declared to be of type `event_t`.
- o. Elements of a struct or union must belong to the same address space. Declaring a struct or union whose elements are in different address spaces is illegal.
- p. Arguments to kernel functions that are declared to be a struct or union do not allow OpenCL objects to be passed as elements of the struct or union.
- q. The type qualifiers `const`, `restrict` and `volatile` as defined by the C99 specification are supported. These qualifiers cannot be used with `image2d_t`, `image3d_t`, `image2d_array_t`, `image1d_t`, `image1d_buffer_t` and `image1d_array_t` types. Types other than pointer types shall not use the `restrict` qualifier.
- r. The event type (`event_t`) cannot be used as the type of a kernel function argument. The event type cannot be used to declare a program scope variable. The event type cannot be used to declare a structure or union field. The event type cannot be used with the `__local`, `__constant` and `__global` address space qualifiers.

6.10 Preprocessor Directives and Macros

The preprocessing directives defined by the C99 specification are supported.

The **# pragma** directive is described as:

```
# pragma pp-tokensopt new-line
```

A **# pragma** directive where the preprocessing token **OPENCL** (used instead of **STDC**) does not immediately follow **pragma** in the directive (prior to any macro replacement) causes the implementation to behave in an implementation-defined manner. The behavior might cause translation to fail or cause the translator or the resulting program to behave in a non-conforming manner. Any such **pragma** that is not recognized by the implementation is ignored. If the preprocessing token **OPENCL** does immediately follow **pragma** in the directive (prior to any macro replacement), then no macro replacement is performed on the directive, and the directive shall have one of the following forms whose meanings are described elsewhere:

```
#pragma OPENCL FP_CONTRACT on-off-switch  
on-off-switch: one of ON OFF DEFAULT
```

```
#pragma OPENCL EXTENSION extensionname : behavior
```

```
#pragma OPENCL EXTENSION all : behavior
```

The following predefined macro names are available.

`__FILE__` The presumed name of the current source file (a character string literal).

`__LINE__` The presumed line number (within the current source file) of the current source line (an integer constant).

`__OPENCL_VERSION__` substitutes an integer reflecting the version number of the OpenCL supported by the OpenCL device. The version of OpenCL described in this document will have `__OPENCL_VERSION__` substitute the integer 120.

`CL_VERSION_1_0` substitutes the integer 100 reflecting the OpenCL 1.0 version.

`CL_VERSION_1_1` substitutes the integer 110 reflecting the OpenCL 1.1 version.

`CL_VERSION_1_2` substitutes the integer 120 reflecting the OpenCL 1.2 version.

`__OPENCL_C_VERSION__` substitutes an integer reflecting the OpenCL C version specified by the `-cl-std` build option (specified in *section 5.6.4.5*) to **clBuildProgram** or **clCompileProgram**. If the `-cl-std` build option is not specified, the OpenCL C version

supported by the compiler for this OpenCL device will be used. The version of OpenCL C described in this document will have `__OPENCL_C_VERSION__` substitute the integer 120.

`__ENDIAN_LITTLE__` is used to determine if the OpenCL device is a little endian architecture or a big endian architecture (an integer constant of 1 if device is little endian and is undefined otherwise). Also refer to `CL_DEVICE_ENDIAN_LITTLE` specified in *table 4.3*.

`__kernel_exec(X, typen)` (and `kernel_exec(X, typen)`) is defined as
`__kernel __attribute__((work_group_size_hint(X, 1, 1))) \`
`__attribute__((vec_type_hint(typen)))`

`__IMAGE_SUPPORT__` is used to determine if the OpenCL device supports images. This is an integer constant of 1 if images are supported and is undefined otherwise. Also refer to `CL_DEVICE_IMAGE_SUPPORT` specified in *table 4.3*.

`__FAST_RELAXED_MATH__` is used to determine if the `-cl-fast-relaxed-math` optimization option is specified in build options given to **clBuildProgram** or **clCompileProgram**. This is an integer constant of 1 if the `-cl-fast-relaxed-math` build option is specified and is undefined otherwise.

The macro names defined by the C99 specification but not currently supported by OpenCL are reserved for future use.

The predefined identifier `__func__` is available.

6.11 Attribute Qualifiers

This section describes the syntax with which `__attribute__` may be used, and the constructs to which attribute specifiers bind.

An attribute specifier is of the form `__attribute__ ((attribute-list))`.

An attribute list is defined as:

```
attribute-list:
    attributeopt
    attribute-list , attributeopt

attribute:
    attribute-token attribute-argument-clauseopt

attribute-token:
    identifier

attribute-argument-clause:
    ( attribute-argument-list )

attribute-argument-list:
    attribute-argument
    attribute-argument-list, attribute-argument

attribute-argument:
    assignment-expression
```

This syntax is taken directly from GCC but unlike GCC, which allows attributes to be applied only to functions, types, and variables, OpenCL attributes can be associated with:

- + types;
- + functions;
- + variables;
- + blocks; and
- + control-flow statements.

In general, the rules for how an attribute binds, for a given context, are non-trivial and the reader is pointed to GCC's documentation and Maurer and Wong's paper [See 16. and 17. in *section 11 – References*] for the details.

6.11.1 Specifying Attributes of Types

The keyword `__attribute__` allows you to specify special attributes of `enum`, `struct` and `union` types when you define such types. This keyword is followed by an attribute specification inside double parentheses. Two attributes are currently defined for types: `aligned`, and `packed`.

You may specify type attributes in an `enum`, `struct` or `union` type declaration or definition, or for other types in a `typedef` declaration.

For an `enum`, `struct` or `union` type, you may specify attributes either between the `enum`, `struct` or `union` tag and the name of the type, or just past the closing curly brace of the *definition*. The former syntax is preferred.

```
aligned (alignment)
```

This attribute specifies a minimum alignment (in bytes) for variables of the specified type. For example, the declarations:

```
struct S { short f[3]; } __attribute__ ((aligned (8)));  
typedef int more_aligned_int __attribute__ ((aligned (8)));
```

force the compiler to insure (as far as it can) that each variable whose type is `struct S` or `more_aligned_int` will be allocated and aligned *at least* on a 8-byte boundary.

Note that the alignment of any given `struct` or `union` type is required by the ISO C standard to be at least a perfect multiple of the lowest common multiple of the alignments of all of the members of the `struct` or `union` in question and must also be a power of two. This means that you *can* effectively adjust the alignment of a `struct` or `union` type by attaching an `aligned` attribute to any one of the members of such a type, but the notation illustrated in the example above is a more obvious, intuitive, and readable way to request the compiler to adjust the alignment of an entire `struct` or `union` type.

As in the preceding example, you can explicitly specify the alignment (in bytes) that you wish the compiler to use for a given `struct` or `union` type. Alternatively, you can leave out the alignment factor and just ask the compiler to align a type to the maximum useful alignment for the target machine you are compiling for. For example, you could write:

```
struct S { short f[3]; } __attribute__ ((aligned));
```

Whenever you leave out the alignment factor in an `aligned` attribute specification, the compiler automatically sets the alignment for the type to the largest alignment which is ever used for any data type on the target machine you are compiling for. In the example

above, the size of each `short` is 2 bytes, and therefore the size of the entire `struct S` type is 6 bytes. The smallest power of two which is greater than or equal to that is 8, so the compiler sets the alignment for the entire `struct S` type to 8 bytes.

Note that the effectiveness of aligned attributes may be limited by inherent limitations of the OpenCL device and compiler. For some devices, the OpenCL compiler may only be able to arrange for variables to be aligned up to a certain maximum alignment. If the OpenCL compiler is only able to align variables up to a maximum of 8 byte alignment, then specifying `aligned(16)` in an `__attribute__` will still only provide you with 8 byte alignment. See your platform-specific documentation for further information.

The `aligned` attribute can only increase the alignment; but you can decrease it by specifying `packed` as well. See below.

`packed`

This attribute, attached to `struct` or `union` type definition, specifies that each member of the structure or union is placed to minimize the memory required. When attached to an `enum` definition, it indicates that the smallest integral type should be used.

Specifying this attribute for `struct` and `union` types is equivalent to specifying the `packed` attribute on each of the structure or union members.

In the following example `struct my_packed_struct`'s members are packed closely together, but the internal layout of its `s` member is not packed. To do that, `struct my_unpacked_struct` would need to be packed, too.

```
struct my_unpacked_struct
{
    char c;
    int i;
};

struct __attribute__((packed)) my_packed_struct
{
    char c;
    int i;
    struct my_unpacked_struct s;
};
```

You may only specify this attribute on the definition of a `enum`, `struct` or `union`, not on a `typedef` which does not also define the enumerated type, structure or union.

6.11.2 Specifying Attributes of Functions

Refer to *section 6.7* for the function attribute qualifiers currently supported.

6.11.3 Specifying Attributes of Variables

The keyword `__attribute__` allows you to specify special attributes of variables or structure fields. This keyword is followed by an attribute specification inside double parentheses. The following attribute qualifiers are currently defined:

`aligned` (*alignment*)

This attribute specifies a minimum alignment for the variable or structure field, measured in bytes. For example, the declaration:

```
int x __attribute__ ((aligned (16))) = 0;
```

causes the compiler to allocate the global variable `x` on a 16-byte boundary. The alignment value specified must be a power of two.

You can also specify the alignment of structure fields. For example, to create double-word aligned `int` pair, you could write:

```
struct foo { int x[2] __attribute__ ((aligned (8))); };
```

This is an alternative to creating a union with a `double` member that forces the union to be double-word aligned.

As in the preceding examples, you can explicitly specify the alignment (in bytes) that you wish the compiler to use for a given variable or structure field. Alternatively, you can leave out the alignment factor and just ask the compiler to align a variable or field to the maximum useful alignment for the target machine you are compiling for. For example, you could write:

```
short array[3] __attribute__ ((aligned));
```

Whenever you leave out the alignment factor in an `aligned` attribute specification, the OpenCL compiler automatically sets the alignment for the declared variable or field to the largest alignment which is ever used for any data type on the target device you are compiling for.

When used on a `struct`, or `struct` member, the `aligned` attribute can only increase

the alignment; in order to decrease it, the `packed` attribute must be specified as well. When used as part of a `typedef`, the `aligned` attribute can both increase and decrease alignment, and specifying the `packed` attribute will generate a warning.

Note that the effectiveness of aligned attributes may be limited by inherent limitations of the OpenCL device and compiler. For some devices, the OpenCL compiler may only be able to arrange for variables to be aligned up to a certain maximum alignment. If the OpenCL compiler is only able to align variables up to a maximum of 8 byte alignment, then specifying `aligned(16)` in an `__attribute__` will still only provide you with 8 byte alignment. See your platform-specific documentation for further information.

`packed`

The `packed` attribute specifies that a variable or structure field should have the smallest possible alignment—one byte for a variable, unless you specify a larger value with the `aligned` attribute.

Here is a structure in which the field `x` is packed, so that it immediately follows `a`:

```
struct foo
{
    char a;
    int x[2] __attribute__((packed));
};
```

An attribute list placed at the beginning of a user-defined type applies to the variable of that type and not the type, while attributes following the type body apply to the type. For example:

```
/* a has alignment of 128 */
__attribute__((aligned(128))) struct A {int i;} a;

/* b has alignment of 16 */
__attribute__((aligned(16))) struct B {double d;}
    __attribute__((aligned(32))) b ;

struct A a1; /* a1 has alignment of 4 */

struct B b1; /* b1 has alignment of 32 */
```

`endian` (*endianness*)

The `endian` attribute determines the byte ordering of a variable. *endianness* can be set to `host` indicating the variable uses the endianness of the host processor or can be set to `device` indicating the variable uses the endianness of the device on which the kernel will be executed. The default is `device`.

For example:

```
global float4 *p __attribute__((endian(host)));
```

specifies that data stored in memory pointed to by `p` will be in the host endian format.

The endian attribute can only be applied to pointer types that are in the `global` or `constant` address space. The `endian` attribute cannot be used for variables that are not a pointer type. The endian attribute value for both pointers must be the same when one pointer is assigned to another.

6.11.4 Specifying Attributes of Blocks and Control-Flow-Statements

For basic blocks and control-flow-statements the attribute is placed before the structure in question, for example:

```
__attribute__((attr1)) {...}  
for __attribute__((attr2)) (...) __attribute__((attr3)) {...}
```

Here `attr1` applies to the block in braces and `attr2` and `attr3` apply to the loop's control construct and body, respectively.

No attribute qualifiers for blocks and control-flow-statements are currently defined.

6.11.5 Extending Attribute Qualifiers

The attribute syntax can be extended for standard language extensions and vendor specific extensions. Any extensions should follow the naming conventions outlined in the introduction to *section 9* in the OpenCL 1.2 Extension Specification.

Attributes are intended as useful hints to the compiler. It is our intention that a particular implementation of OpenCL be free to ignore all attributes and the resulting executable binary will produce the same result. This does not preclude an implementation from making use of the additional information provided by attributes and performing optimizations or other transformations as it sees fit. In this case it is the programmer's responsibility to guarantee that the information provided is in some sense correct.

6.12 Built-in Functions

The OpenCL C programming language provides a rich set of built-in functions for scalar and vector operations. Many of these functions are similar to the function names provided in common C libraries but they support scalar and vector argument types. Applications should use the built-in functions wherever possible instead of writing their own version.

User defined OpenCL C functions, behave per C standard rules for functions (C99, TC2, Section 6.9.1). On entry to the function, the size of each variably modified parameter is evaluated and the value of each argument expression is converted to the type of the corresponding parameter as per usual arithmetic conversion rules described in *section 6.2.6*. Built-in functions described in this section behave similarly, except that in order to avoid ambiguity between multiple forms of the same built-in function, implicit scalar widening shall not occur. Note that some built-in functions described in this section do have forms that operate on mixed scalar and vector types, however.

6.12.1 Work-Item Functions

Table 6.7 describes the list of built-in work-item functions that can be used to query the number of dimensions, the global and local work size specified to **clEnqueueNDRangeKernel**, and the global and local identifier of each work-item when this kernel is being executed on a device. The number of dimensions, the global and local work size when executing a kernel using the function **clEnqueueTask** is one.



| Function | Description |
|--|---|
| <code>uint get_work_dim ()</code> | Returns the number of dimensions in use. This is the value given to the <i>work_dim</i> argument specified in clEnqueueNDRangeKernel . For clEnqueueTask , this returns 1. |
| <code>size_t get_global_size (uint dimindx)</code> | Returns the number of global work-items specified for dimension identified by <i>dimindx</i> . This value is given by the <i>global_work_size</i> argument to clEnqueueNDRangeKernel . Valid values of <i>dimindx</i> are 0 to <code>get_work_dim() - 1</code> . For other values of <i>dimindx</i> , <code>get_global_size()</code> returns 1. For clEnqueueTask , this always returns 1. |
| <code>size_t get_global_id (uint dimindx)</code> | Returns the unique global work-item ID value for dimension identified by <i>dimindx</i> . The global work-item ID specifies the work-item ID based on the number of global work-items specified to execute the kernel. Valid values of <i>dimindx</i> are 0 to <code>get_work_dim() - 1</code> . For |

| | |
|--|---|
| | <p>other values of <i>dimindx</i>, get_global_id() returns 0.</p> <p>For clEnqueueTask, this returns 0.</p> |
| size_t get_local_size (uint <i>dimindx</i>) | <p>Returns the number of local work-items specified in dimension identified by <i>dimindx</i>. This value is given by the <i>local_work_size</i> argument to clEnqueueNDRangeKernel if <i>local_work_size</i> is not NULL; otherwise the OpenCL implementation chooses an appropriate <i>local_work_size</i> value which is returned by this function. Valid values of <i>dimindx</i> are 0 to get_work_dim() – 1. For other values of <i>dimindx</i>, get_local_size() returns 1.</p> <p>For clEnqueueTask, this always returns 1.</p> |
| size_t get_local_id (uint <i>dimindx</i>) | <p>Returns the unique local work-item ID i.e. a work-item within a specific work-group for dimension identified by <i>dimindx</i>. Valid values of <i>dimindx</i> are 0 to get_work_dim() – 1. For other values of <i>dimindx</i>, get_local_id() returns 0.</p> <p>For clEnqueueTask, this returns 0.</p> |
| size_t get_num_groups (uint <i>dimindx</i>) | <p>Returns the number of work-groups that will execute a kernel for dimension identified by <i>dimindx</i>. Valid values of <i>dimindx</i> are 0 to get_work_dim() – 1. For other values of <i>dimindx</i>, get_num_groups () returns 1.</p> <p>For clEnqueueTask, this always returns 1.</p> |
| size_t get_group_id (uint <i>dimindx</i>) | <p>get_group_id returns the work-group ID which is a number from 0 .. get_num_groups(dimindx) – 1. Valid values of <i>dimindx</i> are 0 to get_work_dim() – 1. For other values, get_group_id() returns 0.</p> <p>For clEnqueueTask, this returns 0.</p> |
| size_t get_global_offset (uint <i>dimindx</i>) | <p>get_global_offset returns the offset values specified in <i>global_work_offset</i> argument to clEnqueueNDRangeKernel. Valid values of <i>dimindx</i> are 0 to get_work_dim() – 1. For other values, get_global_offset() returns 0.</p> <p>For clEnqueueTask, this returns 0.</p> |

Table 6.7 *Work-Item Functions Table*

6.12.2 Math Functions

The list of built-in math functions is described in *table 6.8*. The built-in math functions are categorized into the following:

-  A list of built-in functions that have scalar or vector argument versions, and,
-  A list of built-in functions that only take scalar float arguments.

The vector versions of the math functions operate component-wise. The description is per-component.

The built-in math functions are not affected by the prevailing rounding mode in the calling environment, and always return the same value as they would if called with the round to nearest even rounding mode.

Table 6.8 describes the list of built-in math functions that can take scalar or vector arguments. We use the generic type name `gentype` to indicate that the function can take `float`, `float2`, `float3`, `float4`, `float8`, `float16`, `double`, `double2`, `double3`, `double4`, `double8` or `double16` as the type for the arguments. We use the generic type name `gentypef` to indicate that the function can take `float`, `float2`, `float3`, `float4`, `float8`, or `float16` as the type for the arguments. We use the generic type name `gentyped` to indicate that the function can take `double`, `double2`, `double3`, `double4`, `double8` or `double16` as the type for the arguments. For any specific use of a function, the actual type has to be the same for all arguments and the return type, unless otherwise specified.

| Function | Description |
|--|---|
| <code>gentype acos (gentype)</code> | Arc cosine function. |
| <code>gentype acosh (gentype)</code> | Inverse hyperbolic cosine. |
| <code>gentype acospi (gentype x)</code> | Compute $\text{acos}(x) / \pi$. |
| <code>gentype asin (gentype)</code> | Arc sine function. |
| <code>gentype asinh (gentype)</code> | Inverse hyperbolic sine. |
| <code>gentype asinpi (gentype x)</code> | Compute $\text{asin}(x) / \pi$. |
| <code>gentype atan (gentype y_over_x)</code> | Arc tangent function. |
| <code>gentype atan2 (gentype y, gentype x)</code> | Arc tangent of y / x . |
| <code>gentype atanh (gentype)</code> | Hyperbolic arc tangent. |
| <code>gentype atanpi (gentype x)</code> | Compute $\text{atan}(x) / \pi$. |
| <code>gentype atan2pi (gentype y, gentype x)</code> | Compute $\text{atan2}(y, x) / \pi$. |
| <code>gentype cbrt (gentype)</code> | Compute cube-root. |
| <code>gentype ceil (gentype)</code> | Round to integral value using the round to positive infinity rounding mode. |
| <code>gentype copysign (gentype x, gentype y)</code> | Returns x with its sign changed to match the sign of |

| | |
|---|---|
| | y . |
| gentype cos (gentype) | Compute cosine. |
| gentype cosh (gentype) | Compute hyperbolic consine. |
| gentype cospi (gentype x) | Compute cos (πx). |
| gentype erfc (gentype) | Complementary error function. |
| gentype erf (gentype) | Error function encountered in integrating the normal distribution. |
| gentype exp (gentype x) | Compute the base- e exponential of x . |
| gentype exp2 (gentype) | Exponential base 2 function. |
| gentype exp10 (gentype) | Exponential base 10 function. |
| gentype expm1 (gentype x) | Compute $e^x - 1.0$. |
| gentype fabs (gentype) | Compute absolute value of a floating-point number. |
| gentype fdim (gentype x , gentype y) | $x - y$ if $x > y$, $+0$ if x is less than or equal to y . |
| gentype floor (gentype) | Round to integral value using the round to negative infinity rounding mode. |
| gentype fma (gentype a , gentype b , gentype c) | 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. |
| gentype fmax (gentype x , gentype y) gentypef fmax (gentypef x , float y) gentyped fmax (gentyped x , double y) | Returns y if $x < y$, otherwise it returns x . If one argument is a NaN, fmax() returns the other argument. If both arguments are NaNs, fmax() returns a NaN. |
| gentype fmin ⁴⁸ (gentype x , gentype y) gentypef fmin (gentypef x , float y) gentyped fmin (gentyped x , double y) | Returns y if $y < x$, otherwise it returns x . If one argument is a NaN, fmin() returns the other argument. If both arguments are NaNs, fmin() returns a NaN. |
| gentype fmod (gentype x , gentype y) | Modulus. Returns $x - y * \text{trunc}(x/y)$. |
| gentype fract (gentype x , __global gentype $*iptr$) ⁴⁹ gentype fract (gentype x , __local gentype $*iptr$) gentype fract (gentype x , __private gentype $*iptr$) | Returns fmin ($x - \text{floor}(x)$, $0x1.\text{ffffep}-1f$). floor (x) is returned in $iptr$. |
| floatn frexp (floatn x , __global intn $*exp$) floatn frexp (floatn x , __local intn $*exp$) | 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 $* 2^{exp}$. |

⁴⁸ **fmin** and **fmax** behave as defined by C99 and may not match the IEEE 754-2008 definition for **minNum** and **maxNum** with regard to signaling NaNs. Specifically, signaling NaNs may behave as quiet NaNs.

⁴⁹ The **min()** operator is there to prevent **fract**(-small) from returning 1.0. It returns the largest positive floating-point number less than 1.0.

| | |
|--|--|
| floatn frexp (floatn <i>x</i> , __private intn *exp) float frexp (float <i>x</i> , __global int *exp) float frexp (float <i>x</i> , __local int *exp) float frexp (float <i>x</i> , private int *exp) | |
| doublen frexp (doublen <i>x</i> , __global intn *exp) doublen frexp (doublen <i>x</i> , __local intn *exp) doublen frexp (doublen <i>x</i> , __private intn *exp) double frexp (double <i>x</i> , __global int *exp) double frexp (double <i>x</i> , __local int *exp) double frexp (double <i>x</i> , private int *exp) | Extract mantissa and exponent from <i>x</i> . For each component the mantissa returned is a float with magnitude in the interval [1/2, 1) or 0. Each component of <i>x</i> equals mantissa returned * 2 ^{exp} . |
| gentype hypot (gentype <i>x</i> , gentype <i>y</i>) | Compute the value of the square root of $x^2 + y^2$ without undue overflow or underflow. |
| intn ilogb (floatn <i>x</i>) int ilogb (float <i>x</i>) intn ilogb (doublen <i>x</i>) int ilogb (double <i>x</i>) | Return the exponent as an integer value. |
| floatn ldexp (floatn <i>x</i> , intn <i>k</i>) floatn ldexp (floatn <i>x</i> , int <i>k</i>) float ldexp (float <i>x</i> , int <i>k</i>) doublen ldexp (doublen <i>x</i> , intn <i>k</i>) doublen ldexp (doublen <i>x</i> , int <i>k</i>) double ldexp (double <i>x</i> , int <i>k</i>) | Multiply <i>x</i> by 2 to the power <i>k</i> . |
| gentype lgamma (gentype <i>x</i>) floatn lgamma_r (floatn <i>x</i> , __global intn *signp) floatn lgamma_r (floatn <i>x</i> , __local intn *signp) floatn lgamma_r (floatn <i>x</i> , __private intn *signp) float lgamma_r (float <i>x</i> , __global int *signp) float lgamma_r (float <i>x</i> , __local int *signp) float lgamma_r (float <i>x</i> , private int *signp) | 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 <i>signp</i> argument of lgamma_r . |

| | |
|---|--|
| <pre> __private int *signp) double lgamma_r (double x, __global intn *signp) double lgamma_r (double x, __local intn *signp) double lgamma_r (double x, __private intn *signp) double lgamma_r (double x, __global int *signp) double lgamma_r (double x, __local int *signp) double lgamma_r (double x, __private int *signp) </pre> | |
| gentype log (gentype) | Compute natural logarithm. |
| gentype log2 (gentype) | Compute a base 2 logarithm. |
| gentype log10 (gentype) | Compute a base 10 logarithm. |
| gentype log1p (gentype <i>x</i>) | Compute $\log_e(1.0 + x)$. |
| gentype logb (gentype <i>x</i>) | Compute the exponent of <i>x</i> , which is the integral part of $\log_r x $. |
| gentype mad (gentype <i>a</i> , gentype <i>b</i> , gentype <i>c</i>) | mad approximates $a * b + c$. Whether or how the product of $a * b$ is rounded and how supernormal or subnormal intermediate products are handled is not defined. mad is intended to be used where speed is preferred over accuracy ⁵⁰ . |
| gentype maxmag (gentype <i>x</i> , gentype <i>y</i>) | Returns <i>x</i> if $ x > y $, <i>y</i> if $ y > x $, otherwise fmax (<i>x</i> , <i>y</i>). |
| gentype minmag (gentype <i>x</i> , gentype <i>y</i>) | Returns <i>x</i> if $ x < y $, <i>y</i> if $ y < x $, otherwise fmin (<i>x</i> , <i>y</i>). |
| gentype modf (gentype <i>x</i> , __global gentype * <i>iptr</i>) gentype modf (gentype <i>x</i> , __local gentype * <i>iptr</i>) gentype modf (gentype <i>x</i> , __private gentype * <i>iptr</i>) | Decompose a floating-point number. The modf function breaks the argument <i>x</i> 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 <i>iptr</i> . |
| floatn nan (uintn <i>nancode</i>) float nan (uint <i>nancode</i>) doublen nan (ulongn <i>nancode</i>) double nan (ulong <i>nancode</i>) | Returns a quiet NaN. The <i>nancode</i> may be placed in the significand of the resulting NaN. |
| gentype nextafter (gentype <i>x</i> , gentype <i>y</i>) | Computes the next representable single-precision floating-point value following <i>x</i> in the direction of <i>y</i> . Thus, if <i>y</i> is less than <i>x</i> , nextafter () returns the |

⁵⁰ The user is cautioned that for some usages, e.g. **mad**(*a*, *b*, -*a***b*), the definition of **mad**() is loose enough that almost any result is allowed from **mad**() for some values of *a* and *b*.

| | |
|--|---|
| | largest representable floating-point number less than x . |
| gentype pow (gentype x , gentype y) | Compute x to the power y . |
| floatn pown (floatn x , intn y) float pown (float x , int y) doublen pown (doublen x , intn y) double pown (double x , int y) | Compute x to the power y , where y is an integer. |
| gentype powr (gentype x , gentype y) | Compute x to the power y , where x is ≥ 0 . |
| gentype remainder (gentype x , gentype y) | Compute the value r such that $r = x - n*y$, where n is the integer nearest the exact value of x/y . If there are two integers closest to x/y , n shall be the even one. If r is zero, it is given the same sign as x . |
| floatn remquo (floatn x , floatn y , __global intn * <i>quo</i>) floatn remquo (floatn x , floatn y , __local intn * <i>quo</i>) floatn remquo (floatn x , floatn y , __private intn * <i>quo</i>) float remquo (float x , float y , __global int * <i>quo</i>) float remquo (float x , float y , __local int * <i>quo</i>) float remquo (float x , float y , __private int * <i>quo</i>) | The remquo function computes the value r such that $r = x - k*y$, where k is the integer nearest the exact value of x/y . If there are two integers closest to x/y , k shall be the even one. If r is zero, it is given the same sign as x . This is the same value that is returned by the remainder function. remquo 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 <i>quo</i> . |
| doublen remquo (doublen x , doublen y , __global intn * <i>quo</i>) doublen remquo (doublen x , doublen y , __local intn * <i>quo</i>) doublen remquo (doublen x , doublen y , __private intn * <i>quo</i>) double remquo (double x , | The remquo function computes the value r such that $r = x - k*y$, where k is the integer nearest the exact value of x/y . If there are two integers closest to x/y , k shall be the even one. If r is zero, it is given the same sign as x . This is the same value that is returned by the remainder function. remquo 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 <i>quo</i> . |

| | |
|---|---|
| double remquo (double <i>x</i> , double <i>y</i> , __global int * <i>quo</i>) double remquo (double <i>x</i> , double <i>y</i> , __local int * <i>quo</i>) double remquo (double <i>x</i> , double <i>y</i> , private int * <i>quo</i>) | |
| gentype rint (gentype) | Round to integral value (using round to nearest even rounding mode) in floating-point format. Refer to section 7.1 for description of rounding modes. |
| floatn rootn (floatn <i>x</i> , intn <i>y</i>) float rootn (float <i>x</i> , int <i>y</i>) doublen rootn (doublen <i>x</i> , intn <i>y</i>) doublen rootn (double <i>x</i> , int <i>y</i>) | Compute <i>x</i> to the power 1/ <i>y</i> . |
| gentype round (gentype <i>x</i>) | Return the integral value nearest to <i>x</i> rounding halfway cases away from zero, regardless of the current rounding direction. |
| gentype rsqrt (gentype) | Compute inverse square root. |
| gentype sin (gentype) | Compute sine. |
| gentype sincos (gentype <i>x</i> , __global gentype * <i>cosval</i>) gentype sincos (gentype <i>x</i> , __local gentype * <i>cosval</i>) gentype sincos (gentype <i>x</i> , __private gentype * <i>cosval</i>) | Compute sine and cosine of <i>x</i> . The computed sine is the return value and computed cosine is returned in <i>cosval</i> . |
| gentype sinh (gentype) | Compute hyperbolic sine. |
| gentype sinpi (gentype <i>x</i>) | Compute sin (πx). |
| gentype sqrt (gentype) | Compute square root. |
| gentype tan (gentype) | Compute tangent. |
| gentype tanh (gentype) | Compute hyperbolic tangent. |
| gentype tanpi (gentype <i>x</i>) | Compute tan (πx). |
| gentype tgamma (gentype) | Compute the gamma function. |
| gentype trunc (gentype) | Round to integral value using the round to zero rounding mode. |

Table 6.8 *Scalar and Vector Argument Built-in Math Function Table*

Table 6.9 describes the following functions:

- ✚ A subset of functions from *table 6.8* that are defined with the `half_` prefix . These functions are implemented with a minimum of 10-bits of accuracy i.e. an ULP value ≤ 8192 ulp.

- ✚ A subset of functions from *table 6.8* that are defined with the `native_` prefix. These functions may map to one or more native device instructions and will typically have better performance compared to the corresponding functions (without the `native_` prefix) described in *table 6.8*. The accuracy (and in some cases the input range(s)) of these functions is implementation-defined.
- ✚ `half_` and `native_` functions for following basic operations: divide and reciprocal.

We use the generic type name `gentype` to indicate that the functions in *table 6.9* can take `float`, `float2`, `float3`, `float4`, `float8` or `float16` as the type for the arguments.

| Function | Description |
|---|---|
| <code>gentype half_cos</code> (<code>gentype x</code>) | Compute cosine. x must be in the range $-2^{16} \dots +2^{16}$. |
| <code>gentype half_divide</code> (<code>gentype x</code> , <code>gentype y</code>) | Compute x / y . |
| <code>gentype half_exp</code> (<code>gentype x</code>) | Compute the base- e exponential of x . |
| <code>gentype half_exp2</code> (<code>gentype x</code>) | Compute the base- 2 exponential of x . |
| <code>gentype half_exp10</code> (<code>gentype x</code>) | Compute the base- 10 exponential of x . |
| <code>gentype half_log</code> (<code>gentype x</code>) | Compute natural logarithm. |
| <code>gentype half_log2</code> (<code>gentype x</code>) | Compute a base 2 logarithm. |
| <code>gentype half_log10</code> (<code>gentype x</code>) | Compute a base 10 logarithm. |
| <code>gentype half_powr</code> (<code>gentype x</code> , <code>gentype y</code>) | Compute x to the power y , where x is ≥ 0 . |
| <code>gentype half_recip</code> (<code>gentype x</code>) | Compute reciprocal. |
| <code>gentype half_sqrt</code> (<code>gentype x</code>) | Compute inverse square root. |
| <code>gentype half_sin</code> (<code>gentype x</code>) | Compute sine. x must be in the range $-2^{16} \dots +2^{16}$. |
| <code>gentype half_sqrt</code> (<code>gentype x</code>) | Compute square root. |
| <code>gentype half_tan</code> (<code>gentype x</code>) | Compute tangent. x must be in the range $-2^{16} \dots +2^{16}$. |
| <code>gentype native_cos</code> (<code>gentype x</code>) | Compute cosine over an implementation-defined range. The maximum error is implementation-defined. |
| <code>gentype native_divide</code> (<code>gentype x</code> , <code>gentype y</code>) | Compute x / y over an implementation-defined range. The maximum error is implementation-defined. |
| <code>gentype native_exp</code> (<code>gentype x</code>) | Compute the base- e exponential of x over an implementation-defined range. The maximum error is implementation-defined. |
| <code>gentype native_exp2</code> (<code>gentype x</code>) | Compute the base- 2 exponential of x over an implementation-defined range. The maximum error is implementation-defined. |
| <code>gentype native_exp10</code> (<code>gentype x</code>) | Compute the base- 10 exponential of x over an implementation-defined range. The maximum error is implementation-defined. |
| <code>gentype native_log</code> (<code>gentype x</code>) | Compute natural logarithm over an implementation- |

| | |
|--|---|
| | defined range. The maximum error is implementation-defined. |
| gentype native_log2 (gentype <i>x</i>) | Compute a base 2 logarithm over an implementation-defined range. The maximum error is implementation-defined. |
| gentype native_log10 (gentype <i>x</i>) | Compute a base 10 logarithm over an implementation-defined range. The maximum error is implementation-defined. |
| gentype native_powr (gentype <i>x</i> , gentype <i>y</i>) | Compute <i>x</i> to the power <i>y</i> , where <i>x</i> is ≥ 0 . The range of <i>x</i> and <i>y</i> are implementation-defined. The maximum error is implementation-defined. |
| gentype native_recip (gentype <i>x</i>) | Compute reciprocal over an implementation-defined range. The maximum error is implementation-defined. |
| gentype native_rsqrt (gentype <i>x</i>) | Compute inverse square root over an implementation-defined range. The maximum error is implementation-defined. |
| gentype native_sin (gentype <i>x</i>) | Compute sine over an implementation-defined range. The maximum error is implementation-defined. |
| gentype native_sqrt (gentype <i>x</i>) | Compute square root over an implementation-defined range. The maximum error is implementation-defined. |
| gentype native_tan (gentype <i>x</i>) | Compute tangent over an implementation-defined range. The maximum error is implementation-defined. |

Table 6.9 *Scalar and Vector Argument Built-in half_ and native_ Math Functions*

Support for denormal values is optional for **half_** functions. The **half_** functions may return any result allowed by *section 7.5.3*, even when `-cl-denorms-are-zero` (see *section 5.6.4.2*) is not in force. Support for denormal values is implementation-defined for **native_** functions.

The following symbolic constants are available. Their values are of type `float` and are accurate within the precision of a single precision floating-point number.

| Constant Name | Description |
|------------------|--|
| MAXFLOAT | Value of maximum non-infinite single-precision floating-point number. |
| HUGE_VALF | A positive float constant expression. HUGE_VALF evaluates to +infinity. Used as an error value returned by the built-in math functions. |
| INFINITY | A constant expression of type <code>float</code> representing positive or unsigned infinity. |
| NAN | A constant expression of type <code>float</code> representing a quiet NaN. |

If double precision is supported by the device, the following symbolic constant will also be available:

| Constant Name | Description |
|---------------|--|
| HUGE_VAL | A positive double constant expression. HUGE_VAL evaluates to +infinity. Used as an error value returned by the built-in math functions. |

6.12.2.1 Floating-point macros and pragmas

The **FP_CONTRACT** pragma can be used to allow (if the state is `on`) or disallow (if the state is `off`) the implementation to contract expressions. Each pragma can occur either outside external declarations or preceding all explicit declarations and statements inside a compound statement. When outside external declarations, the pragma takes effect from its occurrence until another **FP_CONTRACT** pragma is encountered, or until the end of the translation unit. When inside a compound statement, the pragma takes effect from its occurrence until another **FP_CONTRACT** pragma is encountered (including within a nested compound statement), or until the end of the compound statement; at the end of a compound statement the state for the pragma is restored to its condition just before the compound statement. If this pragma is used in any other context, the behavior is undefined.

The pragma definition to set **FP_CONTRACT** is:

```
#pragma OPENCL FP_CONTRACT on-off-switch
```

on-off-switch is one of:

ON, **OFF** or **DEFAULT**.

The **DEFAULT** value is **ON**.

The **FP_FAST_FMAF** macro indicates whether the **fma** function is fast compared with direct code for single precision floating-point. If defined, the **FP_FAST_FMAF** macro shall indicate that the **fma** function generally executes about as fast as, or faster than, a multiply and an add of **float** operands.

The macro names given in the following list must use the values specified. These constant expressions are suitable for use in `#if` preprocessing directives.

```
#define FLT_DIG 6
#define FLT_MANT_DIG 24
#define FLT_MAX_10_EXP +38
#define FLT_MAX_EXP +128
#define FLT_MIN_10_EXP -37
#define FLT_MIN_EXP -125
#define FLT_RADIX 2
#define FLT_MAX 0x1.fffffep127f
#define FLT_MIN 0x1.0p-126f
#define FLT_EPSILON 0x1.0p-23f
```

The following table describes the built-in macro names given above in the OpenCL C programming language and the corresponding macro names available to the application.

| Macro in OpenCL Language | Macro for application |
|-----------------------------|--------------------------------|
| <code>FLT_DIG</code> | <code>CL_FLT_DIG</code> |
| <code>FLT_MANT_DIG</code> | <code>CL_FLT_MANT_DIG</code> |
| <code>FLT_MAX_10_EXP</code> | <code>CL_FLT_MAX_10_EXP</code> |
| <code>FLT_MAX_EXP</code> | <code>CL_FLT_MAX_EXP</code> |
| <code>FLT_MIN_10_EXP</code> | <code>CL_FLT_MIN_10_EXP</code> |
| <code>FLT_MIN_EXP</code> | <code>CL_FLT_MIN_EXP</code> |
| <code>FLT_RADIX</code> | <code>CL_FLT_RADIX</code> |
| <code>FLT_MAX</code> | <code>CL_FLT_MAX</code> |
| <code>FLT_MIN</code> | <code>CL_FLT_MIN</code> |
| <code>FLT_EPSILON</code> | <code>CL_FLT_EPSILON</code> |

The following macros shall expand to integer constant expressions whose values are returned by `ilogb(x)` if x is zero or NaN, respectively. The value of `FP_ILOGB0` shall be either `{INT_MIN}` or `- {INT_MAX}`. The value of `FP_ILOGBNAN` shall be either `{INT_MAX}` or `{INT_MIN}`.

The following constants are also available. They are of type `float` and are accurate within the precision of the `float` type.

| Constant | Description |
|---------------------------|---------------------------|
| <code>M_E_F</code> | Value of e |
| <code>M_LOG2E_F</code> | Value of $\log_2 e$ |
| <code>M_LOG10E_F</code> | Value of $\log_{10} e$ |
| <code>M_LN2_F</code> | Value of $\log_e 2$ |
| <code>M_LN10_F</code> | Value of $\log_e 10$ |
| <code>M_PI_F</code> | Value of π |
| <code>M_PI_2_F</code> | Value of $\pi / 2$ |
| <code>M_PI_4_F</code> | Value of $\pi / 4$ |
| <code>M_1_PI_F</code> | Value of $1 / \pi$ |
| <code>M_2_PI_F</code> | Value of $2 / \pi$ |
| <code>M_2_SQRTPI_F</code> | Value of $2 / \sqrt{\pi}$ |
| <code>M_SQRT2_F</code> | Value of $\sqrt{2}$ |
| <code>M_SQRT1_2_F</code> | Value of $1 / \sqrt{2}$ |

If double precision is supported by the device, the following macros and constants are also available:

The `FP_FAST_FMA` macro indicates whether the `fma()` family of functions are fast compared with direct code for double precision floating-point. If defined, the `FP_FAST_FMA` macro shall

indicate that the `fma()` function generally executes about as fast as, or faster than, a multiply and an add of `double` operands

The macro names given in the following list must use the values specified. These constant expressions are suitable for use in `#if` preprocessing directives.

```
#define DBL_DIG          15
#define DBL_MANT_DIG    53
#define DBL_MAX_10_EXP  +308
#define DBL_MAX_EXP     +1024
#define DBL_MIN_10_EXP  -307
#define DBL_MIN_EXP     -1021
#define DBL_MAX          0x1.fffffffffffffp1023
#define DBL_MIN          0x1.0p-1022
#define DBL_EPSILON     0x1.0p-52
```

The following table describes the built-in macro names given above in the OpenCL C programming language and the corresponding macro names available to the application.

| Macro in OpenCL Language | Macro for application |
|-----------------------------|--------------------------------|
| <code>DBL_DIG</code> | <code>CL_DBL_DIG</code> |
| <code>DBL_MANT_DIG</code> | <code>CL_DBL_MANT_DIG</code> |
| <code>DBL_MAX_10_EXP</code> | <code>CL_DBL_MAX_10_EXP</code> |
| <code>DBL_MAX_EXP</code> | <code>CL_DBL_MAX_EXP</code> |
| <code>DBL_MIN_10_EXP</code> | <code>CL_DBL_MIN_10_EXP</code> |
| <code>DBL_MIN_EXP</code> | <code>CL_DBL_MIN_EXP</code> |
| <code>DBL_MAX</code> | <code>CL_DBL_MAX</code> |
| <code>DBL_MIN</code> | <code>CL_DBL_MIN</code> |
| <code>DBL_EPSILON</code> | <code>CL_DBL_EPSILON</code> |

The following constants are also available. They are of type `double` and are accurate within the precision of the `double` type.

| Constant | Description |
|-------------------------|---------------------------|
| <code>M_E</code> | Value of e |
| <code>M_LOG2E</code> | Value of $\log_2 e$ |
| <code>M_LOG10E</code> | Value of $\log_{10} e$ |
| <code>M_LN2</code> | Value of $\log_e 2$ |
| <code>M_LN10</code> | Value of $\log_e 10$ |
| <code>M_PI</code> | Value of π |
| <code>M_PI_2</code> | Value of $\pi / 2$ |
| <code>M_PI_4</code> | Value of $\pi / 4$ |
| <code>M_1_PI</code> | Value of $1 / \pi$ |
| <code>M_2_PI</code> | Value of $2 / \pi$ |
| <code>M_2_SQRTPI</code> | Value of $2 / \sqrt{\pi}$ |

| | |
|------------------|-------------------------|
| M_SQRT2 | Value of $\sqrt{2}$ |
| M_SQRT1_2 | Value of $1 / \sqrt{2}$ |

6.12.3 Integer Functions

Table 6.10 describes the built-in integer functions that take scalar or vector arguments. The vector versions of the integer functions operate component-wise. The description is per-component.

We use the generic type name `gentype` to indicate that the function can take `char`, `char{2|3|4|8|16}`, `uchar`, `uchar{2|3|4|8|16}`, `short`, `short{2|3|4|8|16}`, `ushort`, `ushort{2|3|4|8|16}`, `int`, `int{2|3|4|8|16}`, `uint`, `uint{2|3|4|8|16}`, `long`, `long{2|3|4|8|16}`, `ulong`, or `ulong{2|3|4|8|16}` as the type for the arguments. We use the generic type name `ugentype` to refer to unsigned versions of `gentype`. For example, if `gentype` is `char4`, `ugentype` is `uchar4`. We also use the generic type name `sgentype` to indicate that the function can take a scalar data type i.e. `char`, `uchar`, `short`, `ushort`, `int`, `uint`, `long`, or `ulong` as the type for the arguments. For built-in integer functions that take `gentype` and `sgentype` arguments, the `gentype` argument must be a vector or scalar version of the `sgentype` argument. For example, if `sgentype` is `uchar`, `gentype` must be `uchar` or `uchar{2|3|4|8|16}`. For vector versions, `sgentype` is implicitly widened to `gentype` as described in section 6.3.a.

For any specific use of a function, the actual type has to be the same for all arguments and the return type unless otherwise specified.

| Function | Description |
|--|---|
| <code>ugentype abs (gentype x)</code> | Returns $ x $. |
| <code>ugentype abs_diff (gentype x, gentype y)</code> | Returns $ x - y $ without modulo overflow. |
| <code>gentype add_sat (gentype x, gentype y)</code> | Returns $x + y$ and saturates the result. |
| <code>gentype hadd (gentype x, gentype y)</code> | Returns $(x + y) \gg 1$. The intermediate sum does not modulo overflow. |
| <code>gentype rhadd (gentype x, gentype y)⁵¹</code> | Returns $(x + y + 1) \gg 1$. The intermediate sum does not modulo overflow. |
| <code>gentype clamp (gentype x, gentype minval, gentype maxval)</code> <code>gentype clamp (gentype x, sgentype minval, sgentype maxval)</code> | Returns min (max (x , $minval$), $maxval$). Results are undefined if $minval > maxval$. |
| <code>gentype clz (gentype x)</code> | Returns the number of leading 0-bits in x , starting at the most significant bit position. |
| <code>gentype mad_hi (gentype a,</code> | Returns mul_hi (a, b) + c . |

⁵¹ Frequently vector operations need $n + 1$ bits temporarily to calculate a result. The **rhadd** instruction gives you an extra bit without needing to upsample and downsample. This can be a profound performance win.

| | |
|--|---|
| gentype mad_sat (gentype <i>a</i> , gentype <i>b</i> , gentype <i>c</i>) | Returns $a * b + c$ and saturates the result. |
| gentype max (gentype <i>x</i> , gentype <i>y</i>) gentype max (gentype <i>x</i> , sgentype <i>y</i>) | Returns <i>y</i> if $x < y$, otherwise it returns <i>x</i> . |
| gentype min (gentype <i>x</i> , gentype <i>y</i>) gentype min (gentype <i>x</i> , sgentype <i>y</i>) | Returns <i>y</i> if $y < x$, otherwise it returns <i>x</i> . |
| gentype mul_hi (gentype <i>x</i> , gentype <i>y</i>) | Computes $x * y$ and returns the high half of the product of <i>x</i> and <i>y</i> . |
| gentype rotate (gentype <i>v</i> , gentype <i>i</i>) | For each element in <i>v</i> , the bits are shifted left by the number of bits given by the corresponding element in <i>i</i> (subject to usual shift modulo rules described in <i>section 6.3</i>). Bits shifted off the left side of the element are shifted back in from the right. |
| gentype sub_sat (gentype <i>x</i> , gentype <i>y</i>) | Returns $x - y$ and saturates the result. |
| short upsample (char <i>hi</i> , uchar <i>lo</i>) ushort upsample (uchar <i>hi</i> , uchar <i>lo</i>) shortn upsample (charn <i>hi</i> , ucharn <i>lo</i>) ushortn upsample (ucharn <i>hi</i> , ucharn <i>lo</i>) int upsample (short <i>hi</i> , ushort <i>lo</i>) uint upsample (ushort <i>hi</i> , ushort <i>lo</i>) intn upsample (shortn <i>hi</i> , ushortn <i>lo</i>) uintn upsample (ushortn <i>hi</i> , ushortn <i>lo</i>) long upsample (int <i>hi</i> , uint <i>lo</i>) ulong upsample (uint <i>hi</i> , uint <i>lo</i>) longn upsample (intn <i>hi</i> , uintn <i>lo</i>) ulongn upsample (uintn <i>hi</i> , uintn <i>lo</i>) | $result[i] = ((short)hi[i] \ll 8) lo[i]$ $result[i] = ((ushort)hi[i] \ll 8) lo[i]$ $result[i] = ((int)hi[i] \ll 16) lo[i]$ $result[i] = ((uint)hi[i] \ll 16) lo[i]$ $result[i] = ((long)hi[i] \ll 32) lo[i]$ $result[i] = ((ulong)hi[i] \ll 32) lo[i]$ |
| gentype popcount (gentype <i>x</i>) | Returns the number of non-zero bits in <i>x</i> . |

Table 6.10 *Scalar and Vector Integer Argument Built-in Functions*

Table 6.11 describes fast integer functions that can be used for optimizing performance of kernels. We use the generic type name *gentype* to indicate that the function can take *int*, *int2*, *int3*, *int4*, *int8*, *int16*, *uint*, *uint2*, *uint3*, *uint4*, *uint8* or *uint16* as the type for the arguments.

| Function | Description |
|---|--|
| gentype mad24 (gentype <i>x</i> , gentype <i>y</i> , gentype <i>z</i>) | Multiply two 24-bit integer values <i>x</i> and <i>y</i> and add the 32-bit integer result to the 32-bit integer <i>z</i> . Refer to definition of mul24 to see how the 24-bit integer multiplication is performed. |

| | |
|---|---|
| gentype mul24 (gentype <i>x</i> , gentype <i>y</i>) | Multiply two 24-bit integer values <i>x</i> and <i>y</i> . <i>x</i> and <i>y</i> are 32-bit integers but only the low 24-bits are used to perform the multiplication. mul24 should only be used when values in <i>x</i> and <i>y</i> are in the range $[-2^{23}, 2^{23}-1]$ if <i>x</i> and <i>y</i> are signed integers and in the range $[0, 2^{24}-1]$ if <i>x</i> and <i>y</i> are unsigned integers. If <i>x</i> and <i>y</i> are not in this range, the multiplication result is implementation-defined. |
|---|---|

Table 6.11 *Fast Integer Built-in Functions*

The macro names given in the following list must use the values specified. The values shall all be constant expressions suitable for use in `#if` preprocessing directives.

```

#define CHAR_BIT      8
#define CHAR_MAX      SCHAR_MAX
#define CHAR_MIN      SCHAR_MIN
#define INT_MAX        2147483647
#define INT_MIN        (-2147483647 - 1)
#define LONG_MAX       0x7fffffffffffffffL
#define LONG_MIN       (-0x7fffffffffffffffL - 1)
#define SCHAR_MAX     127
#define SCHAR_MIN     (-127 - 1)
#define SHRT_MAX       32767
#define SHRT_MIN       (-32767 - 1)
#define UCHAR_MAX     255
#define USHRT_MAX     65535
#define UINT_MAX       0xffffffff
#define ULONG_MAX     0xffffffffffffffffUL

```

The following table describes the built-in macro names given above in the OpenCL C programming language and the corresponding macro names available to the application.

| Macro in OpenCL Language | Macro for application |
|--------------------------|-----------------------|
| CHAR_BIT | CL_CHAR_BIT |
| CHAR_MAX | CL_CHAR_MAX |
| CHAR_MIN | CL_CHAR_MIN |
| INT_MAX | CL_INT_MAX |
| INT_MIN | CL_INT_MIN |
| LONG_MAX | CL_LONG_MAX |
| LONG_MIN | CL_LONG_MIN |
| SCHAR_MAX | CL_SCHAR_MAX |
| SCHAR_MIN | CL_SCHAR_MIN |
| SHRT_MAX | CL_SHRT_MAX |
| SHRT_MIN | CL_SHRT_MIN |
| UCHAR_MAX | CL_UCHAR_MAX |

| | |
|------------------|---------------------|
| USHRT_MAX | CL_USHRT_MAX |
| UINT_MAX | CL_UINT_MAX |
| ULONG_MAX | CL_ULONG_MAX |

6.12.4 Common Functions⁵²

Table 6.12 describes the list of built-in common functions. These all operate component-wise. The description is per-component. We use the generic type name `gentype` to indicate that the function can take `float`, `float2`, `float3`, `float4`, `float8`, `float16`, `double`, `double2`, `double3`, `double4`, `double8` or `double16` as the type for the arguments. We use the generic type name `gentypef` to indicate that the function can take `float`, `float2`, `float3`, `float4`, `float8`, or `float16` as the type for the arguments. We use the generic type name `gentyped` to indicate that the function can take `double`, `double2`, `double3`, `double4`, `double8` or `double16` as the type for the arguments.

The built-in common functions are implemented using the round to nearest even rounding mode.

| Function | Description |
|---|---|
| <code>gentype clamp</code> (<code>gentype x</code> , <code>gentype minval</code> , <code>gentype maxval</code>) <code>gentypef clamp</code> (<code>gentypef x</code> , <code>float minval</code> , <code>float maxval</code>) <code>gentyped clamp</code> (<code>gentyped x</code> , <code>double minval</code> , <code>double maxval</code>) | Returns <code>fmin(fmax(x, minval), maxval)</code> . Results are undefined if <code>minval > maxval</code> . |
| <code>gentype degrees</code> (<code>gentype radians</code>) | Converts <code>radians</code> to degrees, i.e. $(180 / \pi) * radians$. |
| <code>gentype max</code> (<code>gentype x</code> , <code>gentype y</code>) <code>gentypef max</code> (<code>gentypef x</code> , <code>float y</code>) <code>gentyped max</code> (<code>gentyped x</code> , <code>double y</code>) | Returns <code>y</code> if <code>x < y</code> , otherwise it returns <code>x</code> . If <code>x</code> or <code>y</code> are infinite or NaN, the return values are undefined. |
| <code>gentype min</code> (<code>gentype x</code> , <code>gentype y</code>) <code>gentypef min</code> (<code>gentypef x</code> , <code>float y</code>) <code>gentyped min</code> (<code>gentyped x</code> , <code>double y</code>) | Returns <code>y</code> if <code>y < x</code> , otherwise it returns <code>x</code> . If <code>x</code> or <code>y</code> are infinite or NaN, the return values are undefined. |
| <code>gentype mix</code> (<code>gentype x</code> , <code>gentype y</code> , <code>gentype a</code>) | Returns the linear blend of <code>x</code> & <code>y</code> implemented as: $x + (y - x) * a$ |

⁵² The `mix` and `smoothstep` functions can be implemented using contractions such as `mad` or `fma`.

| | |
|---|--|
| gentypef mix (gentypef <i>x</i> , gentypef <i>y</i> , float <i>a</i>) gentyped mix (gentyped <i>x</i> , gentyped <i>y</i> , double <i>a</i>) | <i>a</i> must be a value in the range 0.0 ... 1.0. If <i>a</i> is not in the range 0.0 ... 1.0, the return values are undefined. |
| gentype radians (gentype <i>degrees</i>) | Converts <i>degrees</i> to radians, i.e. $(\pi / 180) * \textit{degrees}$. |
| gentype step (gentype <i>edge</i> , gentype <i>x</i>) gentypef step (float <i>edge</i> , gentypef <i>x</i>) gentyped step (double <i>edge</i> , gentyped <i>x</i>) | Returns 0.0 if $x < \textit{edge}$, otherwise it returns 1.0. |
| gentype smoothstep (gentype <i>edge0</i> , gentype <i>edge1</i> , gentype <i>x</i>) gentypef smoothstep (float <i>edge0</i> , float <i>edge1</i> , gentypef <i>x</i>) gentyped smoothstep (double <i>edge0</i> , double <i>edge1</i> , gentyped <i>x</i>) | Returns 0.0 if $x \leq \textit{edge0}$ and 1.0 if $x \geq \textit{edge1}$ and performs smooth Hermite interpolation between 0 and 1 when $\textit{edge0} < x < \textit{edge1}$. This is useful in cases where you would want a threshold function with a smooth transition. This is equivalent to: gentype <i>t</i> ; <i>t</i> = clamp $((x - \textit{edge0}) / (\textit{edge1} - \textit{edge0}), 0, 1)$; return $t * t * (3 - 2 * t)$; Results are undefined if $\textit{edge0} \geq \textit{edge1}$ or if <i>x</i> , <i>edge0</i> or <i>edge1</i> is a NaN. |
| gentype sign (gentype <i>x</i>) | 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 <i>x</i> is a NaN. |

Table 6.12 *Scalar and Vector Argument Built-in Common Function Table*

6.12.5 Geometric Functions⁵³

Table 6.13 describes the list of built-in geometric functions. These all operate component-wise. The description is per-component. `floatn` is `float`, `float2`, `float3`, or `float4` and `doublen` is `double`, `double2`, `double3`, or `double4`. The built-in geometric functions are implemented using the round to nearest even rounding mode.

| 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 <code>float4</code> result returned will be 0.0. |
| <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 fast_length ($p0 - p1$). |
| <code>float fast_length (floatn p)</code> | Returns the length of vector p computed as: half_sqrt ($p.x^2 + p.y^2 + \dots$) |
| <code>floatn fast_normalize (floatn p)</code> | Returns a vector in the same direction as p but with a length of 1. fast_normalize is computed as: $p * \mathbf{half_rsqrt}(p.x^2 + p.y^2 + \dots)$ <p>The result shall be within 8192 ulps error from the infinitely precise result of</p> |

⁵³ The geometric functions can be implemented using contractions such as **mad** or **fma**.

| | |
|--|---|
| | <pre> if (all($p == 0.0f$)) $result = p$; else $result = p / \mathbf{sqrt}(p.x^2 + p.y^2 + \dots)$; </pre> <p>with the following exceptions:</p> <ol style="list-style-type: none"> 1) If the sum of squares is greater than FLT_MAX then the value of the floating-point values in the result vector are undefined. 2) If the sum of squares is less than FLT_MIN then the implementation may return back p. 3) If the device is in “denorms are flushed to zero” mode, individual operand elements with magnitude less than $\mathbf{sqrt}(FLT_MIN)$ may be flushed to zero before proceeding with the calculation. |
|--|---|

Table 6.13 *Scalar and Vector Argument Built-in Geometric Function Table*

6.12.6 Relational Functions

The relational and equality operators (<, <=, >, >=, !=, ==) can be used with scalar and vector built-in types and produce a scalar or vector signed integer result respectively as described in *section 6.3*.

The functions⁵⁴ described in *table 6.14* can be used with built-in scalar or vector types as arguments and return a scalar or vector integer result. The argument type *gentype* refers to the following built-in types: *char*, *charn*, *uchar*, *ucharn*, *short*, *shortn*, *ushort*, *ushortn*, *int*, *intn*, *uint*, *uintn*, *long*, *longn*, *ulong*, *ulongn*, *float*, *floatn*, *double*, and *doublen*. The argument type *igentype* refers to the built-in signed integer types i.e. *char*, *charn*, *short*, *shortn*, *int*, *intn*, *long* and *longn*. The argument type *ugentype* refers to the built-in unsigned integer types i.e. *uchar*, *ucharn*, *ushort*, *ushortn*, *uint*, *uintn*, *ulong* and *ulongn*. *n* is 2, 3, 4, 8, or 16.

The functions **isequal**, **isnotequal**, **isgreater**, **isgreaterequal**, **isless**, **islessequal**, **islessgreater**, **isfinite**, **isinf**, **isnan**, **isnormal**, **isordered**, **isunordered** and **signbit** described in *table 6.14* shall return a 0 if the specified relation is *false* and a 1 if the specified relation is true for scalar argument types. These functions shall return a 0 if the specified relation is *false* and a -1 (i.e. all bits set) if the specified relation is *true* for vector argument types.

The relational functions **isequal**, **isgreater**, **isgreaterequal**, **isless**, **islessequal**, and **islessgreater** always return 0 if either argument is not a number (NaN). **isnotequal** returns 1 if one or both arguments are not a number (NaN) and the argument type is a scalar and returns -1 if one or both arguments are not a number (NaN) and the argument type is a vector.

| Function | Description |
|--|--|
| int isequal (float <i>x</i> , float <i>y</i>) intn isequal (floatn <i>x</i> , floatn <i>y</i>) | Returns the component-wise compare of $x == y$. |
| int isequal (double <i>x</i> , double <i>y</i>) longn isequal (doublen <i>x</i> , doublen <i>y</i>) | |
| int isnotequal (float <i>x</i> , float <i>y</i>) intn isnotequal (floatn <i>x</i> , floatn <i>y</i>) | Returns the component-wise compare of $x != y$. |
| int isnotequal (double <i>x</i> , double <i>y</i>) longn isnotequal (doublen <i>x</i> , doublen <i>y</i>) | |
| int isgreater (float <i>x</i> , float <i>y</i>) | Returns the component-wise compare of $x > y$. |

⁵⁴ If an implementation extends this specification to support IEEE-754 flags or exceptions, then all builtin functions defined in *table 6.14* shall proceed without raising the *invalid* floating-point exception when one or more of the operands are NaNs.

| | |
|--|---|
| <p><code>intn isgreater (floatn x, floatn y)</code></p> <p><code>int isgreater (double x, double y)</code> <code>longn isgreater (doublen x, doublen y)</code></p> | |
| <p><code>int isgreaterequal (float x, float y)</code> <code>intn isgreaterequal (floatn x, floatn y)</code></p> <p><code>int isgreaterequal (double x, double y)</code> <code>longn isgreaterequal (doublen x, doublen y)</code></p> | Returns the component-wise compare of $x \geq y$. |
| <p><code>int isless (float x, float y)</code> <code>intn isless (floatn x, floatn y)</code></p> <p><code>int isless (double x, double y)</code> <code>longn isless (doublen x, doublen y)</code></p> | Returns the component-wise compare of $x < y$. |
| <p><code>int islessequal (float x, float y)</code> <code>intn islessequal (floatn x, floatn y)</code></p> <p><code>int islessequal (double x, double y)</code> <code>longn islessequal (doublen x, doublen y)</code></p> | Returns the component-wise compare of $x \leq y$. |
| <p><code>int islessgreater (float x, float y)</code> <code>intn islessgreater (floatn x, floatn y)</code></p> <p><code>int islessgreater (double x, double y)</code> <code>longn islessgreater (doublen x, doublen y)</code></p> | Returns the component-wise compare of $(x < y) \parallel (x > y)$. |
| | |
| <p><code>int isfinite (float)</code> <code>intn isfinite (floatn)</code></p> <p><code>int isfinite (double)</code> <code>longn isfinite (doublen)</code></p> | Test for finite value. |
| <p><code>int isinf (float)</code> <code>intn isinf (floatn)</code></p> <p><code>int isinf (double)</code> <code>longn isinf (doublen)</code></p> | Test for infinity value (positive or negative). |
| <p><code>int isnan (float)</code> <code>intn isnan (floatn)</code></p> <p><code>int isnan (double)</code> <code>longn isnan (doublen)</code></p> | Test for a NaN. |
| <p><code>int isnormal (float)</code> <code>intn isnormal (floatn)</code></p> <p><code>int isnormal (double)</code></p> | Test for a normal value. |

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

Table 6.14 *Scalar and Vector Relational Functions*

6.12.7 Vector Data Load and Store Functions

Table 6.15 describes the list of supported functions that allow you to read and write vector types from a pointer to memory. We use the generic type `gentype` to indicate the built-in data types `char`, `uchar`, `short`, `ushort`, `int`, `uint`, `long`, `ulong`, `float` or `double`. We use the generic type name `gentypen` to represent `n`-element vectors of `gentype` elements. We use the type name `halfn` to represent `n`-element vectors of `half` elements⁵⁵. The suffix `n` is also used in the function names (i.e. `vloadn`, `vstoren` etc.), where `n` = 2, 3, 4, 8 or 16.

| Function | Description |
|--|---|
| <code>gentypen vloadn</code> (<code>size_t offset</code> , <code>const __global gentype *p</code>) <code>gentypen vloadn</code> (<code>size_t offset</code> , <code>const __local gentype *p</code>) <code>gentypen vloadn</code> (<code>size_t offset</code> , <code>const __constant gentype *p</code>) <code>gentypen vloadn</code> (<code>size_t offset</code> , <code>const __private gentype *p</code>) | Return <code>sizeof (gentypen)</code> bytes of data read from address (<code>p + (offset * n)</code>). The address computed as (<code>p + (offset * n)</code>) must be 8-bit aligned if <code>gentype</code> is <code>char</code> , <code>uchar</code> ; 16-bit aligned if <code>gentype</code> is <code>short</code> , <code>ushort</code> ; 32-bit aligned if <code>gentype</code> is <code>int</code> , <code>uint</code> , <code>float</code> ; 64-bit aligned if <code>gentype</code> is <code>long</code> , <code>ulong</code> . |
| <code>void vstoren</code> (<code>gentypen data</code> , <code>size_t offset</code> , <code>__global gentype *p</code>) <code>void vstoren</code> (<code>gentypen data</code> , <code>size_t offset</code> , <code>__local gentype *p</code>) <code>void vstoren</code> (<code>gentypen data</code> , <code>size_t offset</code> , <code>__private gentype *p</code>) | Write <code>sizeof (gentypen)</code> bytes given by <code>data</code> to address (<code>p + (offset * n)</code>). The address computed as (<code>p + (offset * n)</code>) must be 8-bit aligned if <code>gentype</code> is <code>char</code> , <code>uchar</code> ; 16-bit aligned if <code>gentype</code> is <code>short</code> , <code>ushort</code> ; 32-bit aligned if <code>gentype</code> is <code>int</code> , <code>uint</code> , <code>float</code> ; 64-bit aligned if <code>gentype</code> is <code>long</code> , <code>ulong</code> . |
| <code>float vload_half</code> (<code>size_t offset</code> , <code>const __global half *p</code>) <code>float vload_half</code> (<code>size_t offset</code> , <code>const __local half *p</code>) <code>float vload_half</code> (<code>size_t offset</code> , <code>const __private half *p</code>) | Read <code>sizeof (half)</code> bytes of data from address (<code>p + offset</code>). The data read is interpreted as a half value. The half value is converted to a float value and the float value is returned. The read address computed as (<code>p + offset</code>) must be 16-bit aligned. |

⁵⁵ The `halfn` type is only defined by the `cl_khr_fp16` extension described in section 9.5 of the OpenCL 1.2 Extension Specification.

| | |
|---|--|
| <p style="text-align: center;">const __constant half *p)</p> <p>float vload_half (size_t <i>offset</i>, const __private half *p)</p> | |
| <p>floatn vload_halfn (size_t <i>offset</i>, const __global half *p)</p> <p>floatn vload_halfn (size_t <i>offset</i>, const __local half *p)</p> <p>floatn vload_halfn (size_t <i>offset</i>, const __constant half *p)</p> <p>floatn vload_halfn (size_t <i>offset</i>, const __private half *p)</p> | <p>Read sizeof (halfn) bytes of data from address ($p + (offset * n)$). The data read is interpreted as a halfn value. The halfn value read is converted to a floatn value and the floatn value is returned. The read address computed as ($p + (offset * n)$) must be 16-bit aligned.</p> |
| <p>void vstore_half (float <i>data</i>, size_t <i>offset</i>, __global half *p)</p> <p>void vstore_half_rte (float <i>data</i>, size_t <i>offset</i>, __global half *p)</p> <p>void vstore_half_rtz (float <i>data</i>, size_t <i>offset</i>, __global half *p)</p> <p>void vstore_half_rtp (float <i>data</i>, size_t <i>offset</i>, __global half *p)</p> <p>void vstore_half_rtn (float <i>data</i>, size_t <i>offset</i>, __global half *p)</p> <p>void vstore_half (float <i>data</i>, size_t <i>offset</i>, __local half *p)</p> <p>void vstore_half_rte (float <i>data</i>, size_t <i>offset</i>, __local half *p)</p> <p>void vstore_half_rtz (float <i>data</i>, size_t <i>offset</i>, __local half *p)</p> <p>void vstore_half_rtp (float <i>data</i>, size_t <i>offset</i>, __local half *p)</p> <p>void vstore_half_rtn (float <i>data</i>, size_t <i>offset</i>, __local half *p)</p> <p>void vstore_half (float <i>data</i>, size_t <i>offset</i>, __private half *p)</p> <p>void vstore_half_rte (float <i>data</i>, size_t <i>offset</i>, __private half *p)</p> <p>void vstore_half_rtz (float <i>data</i>, size_t <i>offset</i>, __private half *p)</p> <p>void vstore_half_rtp (float <i>data</i>,</p> | <p>The float value given by <i>data</i> is first converted to a half value using the appropriate rounding mode. The half value is then written to address computed as ($p + offset$). The address computed as ($p + offset$) must be 16-bit aligned.</p> <p>vstore_half uses the default rounding mode. The default rounding mode is round to nearest even.</p> |

| | |
|---|--|
| <pre> size_t offset, __private half *p) void vstore_half_rtn (float data, size_t offset, __private half *p) </pre> | |
| <pre> void vstore_halfn (floatn data, size_t offset, __global half *p) void vstore_halfn_rte (floatn data, size_t offset, __global half *p) void vstore_halfn_rtz (floatn data, size_t offset, __global half *p) void vstore_halfn_rtp (floatn data, size_t offset, __global half *p) void vstore_halfn_rtn (floatn data, size_t offset, __global half *p) void vstore_halfn (floatn data, size_t offset, __local half *p) void vstore_halfn_rte (floatn data, size_t offset, __local half *p) void vstore_halfn_rtz (floatn data, size_t offset, __local half *p) void vstore_halfn_rtp (floatn data, size_t offset, __local half *p) void vstore_halfn_rtn (floatn data, size_t offset, __local half *p) void vstore_halfn (floatn data, size_t offset, __private half *p) void vstore_halfn_rte (floatn data, size_t offset, __private half *p) void vstore_halfn_rtz (floatn data, size_t offset, __private half *p) void vstore_halfn_rtp (floatn data, size_t offset, __private half *p) void vstore_halfn_rtn (floatn data, size_t offset, __private half *p) </pre> | <p>The floatn value given by <i>data</i> is converted to a halfn value using the appropriate rounding mode. The halfn value is then written to address computed as $(p + (offset * n))$. The address computed as $(p + (offset * n))$ must be 16-bit aligned.</p> <p>vstore_halfn uses the default rounding mode. The default rounding mode is round to nearest even.</p> |
| <pre> void vstore_half (double data, size_t offset, __global half *p) void vstore_half_rte (double data, size_t offset, __global half *p) void vstore_half_rtz (double data, size_t offset, __global half *p) void vstore_half_rtp (double data, size_t offset, __global half *p) void vstore_half_rtn (double data, size_t offset, __global half *p) </pre> | <p>The double value given by <i>data</i> is first converted to a half value using the appropriate rounding mode. The half value is then written to address computed as $(p + offset)$. The address computed as $(p + offset)$ must be 16-bit aligned.</p> <p>vstore_half use the default rounding mode. The default rounding mode is round to</p> |

| | |
|---|--|
| <pre> size_t <i>offset</i>, __global half *<i>p</i>) void vstore_half (double <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) void vstore_half_rte (double <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) void vstore_half_rtz (double <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) void vstore_half_rtp (double <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) void vstore_half_rtn (double <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) void vstore_half (double <i>data</i>, size_t <i>offset</i>, __private half *<i>p</i>) void vstore_half_rte (double <i>data</i>, size_t <i>offset</i>, __private half *<i>p</i>) void vstore_half_rtz (double <i>data</i>, size_t <i>offset</i>, __private half *<i>p</i>) void vstore_half_rtp (double <i>data</i>, size_t <i>offset</i>, __private half *<i>p</i>) void vstore_half_rtn (double <i>data</i>, size_t <i>offset</i>, __private half *<i>p</i>) </pre> | <p>nearest even.</p> |
| <pre> void vstore_halfn (doublen <i>data</i>, size_t <i>offset</i>, __global half *<i>p</i>) void vstore_halfn_rte (doublen <i>data</i>, size_t <i>offset</i>, __global half *<i>p</i>) void vstore_halfn_rtz (doublen <i>data</i>, size_t <i>offset</i>, __global half *<i>p</i>) void vstore_halfn_rtp (doublen <i>data</i>, size_t <i>offset</i>, __global half *<i>p</i>) void vstore_halfn_rtn (doublen <i>data</i>, size_t <i>offset</i>, __global half *<i>p</i>) void vstore_halfn (doublen <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) void vstore_halfn_rte (doublen <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) void vstore_halfn_rtz (doublen <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) void vstore_halfn_rtp (doublen <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) void vstore_halfn_rtn (doublen <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) </pre> | <p>The <i>doublen</i> value given by <i>data</i> is converted to a <i>halfn</i> value using the appropriate rounding mode. The <i>halfn</i> value is then written to address computed as $(p + (offset * n))$. The address computed as $(p + (offset * n))$ must be 16-bit aligned.</p> <p>vstore_halfn uses the default rounding mode. The default rounding mode is round to nearest even.</p> |

| | |
|--|---|
| <pre> void vstore_halfn (doublen <i>data</i>, size_t <i>offset</i>, __private half *<i>p</i>) void vstore_halfn_rte (doublen <i>data</i>, size_t <i>offset</i>, __private half *<i>p</i>) void vstore_halfn_rtz (doublen <i>data</i>, size_t <i>offset</i>, __private half *<i>p</i>) void vstore_halfn_rtp (doublen <i>data</i>, size_t <i>offset</i>, __private half *<i>p</i>) void vstore_halfn_rtn (doublen <i>data</i>, size_t <i>offset</i>, __private half *<i>p</i>) </pre> | |
| <pre> floatn vloada_halfn (size_t <i>offset</i>, const __global half *<i>p</i>) floatn vloada_halfn (size_t <i>offset</i>, const __local half *<i>p</i>) floatn vloada_halfn (size_t <i>offset</i>, const __constant half *<i>p</i>) floatn vloada_halfn (size_t <i>offset</i>, const __private half *<i>p</i>) </pre> | <p>For $n = 2, 4, 8$ and 16 read <code>sizeof (halfn)</code> bytes of data from address $(p + (offset * n))$. The data read is interpreted as a <code>halfn</code> value. The <code>halfn</code> value read is converted to a <code>floatn</code> value and the <code>floatn</code> value is returned.</p> <p>The address computed as $(p + (offset * n))$ must be aligned to <code>sizeof (halfn)</code> bytes.</p> <p>For $n = 3$, vloada_half3 reads a <code>half3</code> from address $(p + (offset * 4))$ and returns a <code>float3</code>. The address computed as $(p + (offset * 4))$ must be aligned to <code>sizeof (half) * 4</code> bytes.</p> |
| <pre> void vstorea_halfn (floatn <i>data</i>, size_t <i>offset</i>, __global half *<i>p</i>) void vstorea_halfn_rte (floatn <i>data</i>, size_t <i>offset</i>, __global half *<i>p</i>) void vstorea_halfn_rtz (floatn <i>data</i>, size_t <i>offset</i>, __global half *<i>p</i>) void vstorea_halfn_rtp (floatn <i>data</i>, size_t <i>offset</i>, __global half *<i>p</i>) void vstorea_halfn_rtn (floatn <i>data</i>, size_t <i>offset</i>, __global half *<i>p</i>) void vstorea_halfn (floatn <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) void vstorea_halfn_rte (floatn <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) void vstorea_halfn_rtz (floatn <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) void vstorea_halfn_rtp (floatn <i>data</i>, size_t <i>offset</i>, __local half *<i>p</i>) </pre> | <p>The <code>floatn</code> value given by <i>data</i> is converted to a <code>halfn</code> value using the appropriate rounding mode.</p> <p>For $n = 2, 4, 8$ and 16, the <code>halfn</code> value is written to the address computed as $(p + (offset * n))$. The address computed as $(p + (offset * n))$ must be aligned to <code>sizeof (halfn)</code> bytes.</p> <p>For $n = 3$, the <code>half3</code> value is written to the address computed as $(p + (offset * 4))$. The address computed as $(p + (offset * 4))$ must be aligned to <code>sizeof (half) * 4</code> bytes.</p> <p>vstorea_halfn uses the default rounding mode. The default rounding mode is <code>round to nearest even</code>.</p> |

| | |
|--|---|
| <pre> size_t offset, __local half *p) void vstorea_halfn_rtn (floatn data, size_t offset, __local half *p) void vstorea_halfn (floatn data, size_t offset, __private half *p) void vstorea_halfn_rte (floatn data, size_t offset, __private half *p) void vstorea_halfn_rtz (floatn data, size_t offset, __private half *p) void vstorea_halfn_rtp (floatn data, size_t offset, __private half *p) void vstorea_halfn_rtn (floatn data, size_t offset, __private half *p) </pre> | |
| <pre> void vstorea_halfn (doublen data, size_t offset, __global half *p) void vstorea_halfn_rte (doublen data, size_t offset, __global half *p) void vstorea_halfn_rtz (doublen data, size_t offset, __global half *p) void vstorea_halfn_rtp (doublen data, size_t offset, __global half *p) void vstorea_halfn_rtn (doublen data, size_t offset, __global half *p) void vstorea_halfn (doublen data, size_t offset, __local half *p) void vstorea_halfn_rte (doublen data, size_t offset, __local half *p) void vstorea_halfn_rtz (doublen data, size_t offset, __local half *p) void vstorea_halfn_rtp (doublen data, size_t offset, __local half *p) void vstorea_halfn_rtn (doublen data, size_t offset, __local half *p) void vstorea_halfn (doublen data, size_t offset, __private half *p) void vstorea_halfn_rte (doublen data, size_t offset, __private half *p) void vstorea_halfn_rtz (doublen data, size_t offset, __private half *p) void vstorea_halfn_rtp (doublen data, size_t offset, __private half *p) void vstorea_halfn_rtn (doublen data, size_t offset, __private half *p) </pre> | <p>The <i>doublen</i> value is converted to a <i>halfn</i> value using the appropriate rounding mode.</p> <p>For $n = 2, 4, 8$ or 16, the <i>halfn</i> value is written to the address computed as $(p + (offset * n))$. The address computed as $(p + (offset * n))$ must be aligned to <code>sizeof (halfn)</code> bytes.</p> <p>For $n = 3$, the <i>half3</i> value is written to the address computed as $(p + (offset * 4))$. The address computed as $(p + (offset * 4))$ must be aligned to <code>sizeof (half) * 4</code> bytes.</p> <p>vstorea_halfn uses the default rounding mode. The default rounding mode is round to nearest even.</p> |

| | |
|-----------------------------------|--|
| size_t offset, __private half *p) | |
|-----------------------------------|--|

Table 6.15 *Vector Data Load and Store Functions*⁵⁶

The results of vector data load and store functions are undefined if the address being read from or written to is not correctly aligned as described in *table 6.15*. The pointer argument *p* can be a pointer to `__global`, `__local` or `__private` memory for store functions described in *table 6.15*. The pointer argument *p* can be a pointer to `__global`, `__local`, `__constant` or `__private` memory for load functions described in *table 6.15*.

⁵⁶ **vload3**, and **vload_half3** read *x*, *y*, *z* components from address ($p + (offset * 3)$) into a 3-component vector. **vstore3**, and **vstore_half3** write *x*, *y*, *z* components from a 3-component vector to address ($p + (offset * 3)$).

In addition **vloada_half3** reads *x*, *y*, *z* components from address ($p + (offset * 4)$) into a 3-component vector and **vstorea_half3** writes *x*, *y*, *z* components from a 3-component vector to address ($p + (offset * 4)$).

6.12.8 Synchronization Functions

The OpenCL C programming language implements the following synchronization function.

| Function | Description |
|--|---|
| void barrier (cl_mem_fence_flags <i>flags</i>) | <p>All work-items in a work-group executing the kernel on a processor must execute this function before any are allowed to continue execution beyond the barrier. This function must be encountered by all work-items in a work-group executing the kernel.</p> <p>If barrier is inside a conditional statement, then all work-items must enter the conditional if any work-item enters the conditional statement and executes the barrier.</p> <p>If barrier is inside a loop, all work-items must execute the barrier for each iteration of the loop before any are allowed to continue execution beyond the barrier.</p> <p>The barrier function also queues a memory fence (reads and writes) to ensure correct ordering of memory operations to local or global memory.</p> <p>The <i>flags</i> argument specifies the memory address space and can be set to a combination of the following literal values.</p> <p>CLK_LOCAL_MEM_FENCE - The barrier function will either flush any variables stored in local memory or queue a memory fence to ensure correct ordering of memory operations to local memory.</p> <p>CLK_GLOBAL_MEM_FENCE – The barrier function will queue a memory fence to ensure correct ordering of memory operations to global memory. This can be useful when work-items, for example, write to buffer or image objects and then want to read the updated data.</p> |

Table 6.16 Built-in Synchronization Functions

6.12.9 Explicit Memory Fence Functions

The OpenCL C programming language implements the following explicit memory fence functions to provide ordering between memory operations of a work-item.

| Function | Description |
|--|---|
| void mem_fence (cl_mem_fence_flags <i>flags</i>) | <p>Orders loads and stores of a work-item executing a kernel. This means that loads and stores preceding the mem_fence will be committed to memory before any loads and stores following the mem_fence.</p> <p>The <i>flags</i> argument specifies the memory address space and can be set to a combination of the following literal values:</p> <p>CLK_LOCAL_MEM_FENCE CLK_GLOBAL_MEM_FENCE.</p> |
| void read_mem_fence (cl_mem_fence_flags <i>flags</i>) | <p>Read memory barrier that orders only loads.</p> <p>The <i>flags</i> argument specifies the memory address space and can be set to to a combination of the following literal values:</p> <p>CLK_LOCAL_MEM_FENCE CLK_GLOBAL_MEM_FENCE.</p> |
| void write_mem_fence (cl_mem_fence_flags <i>flags</i>) | <p>Write memory barrier that orders only stores.</p> <p>The <i>flags</i> argument specifies the memory address space and can be set to to a combination of the following literal values:</p> <p>CLK_LOCAL_MEM_FENCE CLK_GLOBAL_MEM_FENCE.</p> |

Table 6.17 Built-in Explicit Memory Fence Functions

6.12.10 Async Copies from Global to Local Memory, Local to Global Memory, and Prefetch

The OpenCL C programming language implements the following functions that provide asynchronous copies between global and local memory and a prefetch from global memory.

We use the generic type name `gentype` to indicate the built-in data types `char`, `char{2|357|4|8|16}`, `uchar`, `uchar{2|3|4|8|16}`, `short`, `short{2|3|4|8|16}`, `ushort`, `ushort{2|3|4|8|16}`, `int`, `int{2|3|4|8|16}`, `uint`, `uint{2|3|4|8|16}`, `long`, `long{2|3|4|8|16}`, `ulong`, `ulong{2|3|4|8|16}`, `float`, `float{2|3|4|8|16}`, or `double`, `double{2|3|4|8|16}` as the type for the arguments unless otherwise stated.

| Function | Description |
|---|---|
| <pre> event_t async_work_group_copy (__local gentype *dst, const __global gentype *src, size_t num_gentypes, event_t event) event_t async_work_group_copy (__global gentype *dst, const __local gentype *src, size_t num_gentypes, event_t event) </pre> | <p>Perform an async copy of <i>num_gentypes</i> <i>gentype</i> elements from <i>src</i> to <i>dst</i>. The async copy is performed by all work-items in a work-group and this built-in function must therefore be encountered by all work-items in a work-group executing the kernel with the same argument values; otherwise the results are undefined.</p> <p>Returns an event object that can be used by wait_group_events to wait for the async copy to finish. The <i>event</i> argument can also be used to associate the async_work_group_copy with a previous async copy allowing an event to be shared by multiple async copies; otherwise <i>event</i> should be zero.</p> <p>If <i>event</i> argument is non-zero, the event object supplied in <i>event</i> argument will be returned.</p> <p>This function does not perform any implicit synchronization of source data such as using a barrier before performing the copy.</p> |

⁵⁷ **async_work_group_copy** and **async_work_group_strided_copy** for 3-component vector types behave as **async_work_group_copy** and **async_work_group_strided_copy** respectively for 4-component vector types.

| | |
|---|--|
| <pre> event_t async_work_group_strided_copy (__local gentype *dst, const __global gentype *src, size_t num_gentypes, size_t src_stride, event_t event) event_t async_work_group_strided_copy (__global gentype *dst, const __local gentype *src, size_t num_gentypes, size_t dst_stride, event_t event) </pre> | <p>Perform an async gather of <i>num_gentypes</i> gentype elements from <i>src</i> to <i>dst</i>. The <i>src_stride</i> is the stride in elements for each gentype element read from <i>src</i>. The <i>dst_stride</i> is the stride in elements for each gentype element written to <i>dst</i>. The async gather is performed by all work-items in a work-group. This built-in function must therefore be encountered by all work-items in a work-group executing the kernel with the same argument values; otherwise the results are undefined.</p> <p>Returns an event object that can be used by wait_group_events to wait for the async copy to finish. The <i>event</i> argument can also be used to associate the async_work_group_strided_copy with a previous async copy allowing an event to be shared by multiple async copies; otherwise <i>event</i> should be zero.</p> <p>If <i>event</i> argument is non-zero, the event object supplied in <i>event</i> argument will be returned.</p> <p>This function does not perform any implicit synchronization of source data such as using a barrier before performing the copy.</p> <p>The behavior of async_work_group_strided_copy is undefined if <i>src_stride</i> or <i>dst_stride</i> is 0, or if the <i>src_stride</i> or <i>dst_stride</i> values cause the <i>src</i> or <i>dst</i> pointers to exceed the upper bounds of the address space during the copy.</p> |
| <pre> void wait_group_events (int num_events, event_t *event_list) </pre> | <p>Wait for events that identify the async_work_group_copy operations to complete. The event objects specified in <i>event_list</i> will be released after the wait is performed.</p> <p>This function must be encountered by all</p> |

| | |
|--|--|
| | work-items in a work-group executing the kernel with the same <i>num_events</i> and event objects specified in <i>event_list</i> ; otherwise the results are undefined. |
| | |
| void prefetch (const __global gentype *p, size_t num_gentypes) | Prefetch <i>num_gentypes</i> * sizeof(<i>gentype</i>) bytes into the global cache. The prefetch instruction is applied to a work-item in a work-group and does not affect the functional behavior of the kernel. |

Table 6.18 *Built-in Async Copy and Prefetch Functions*

NOTE: The kernel must wait for the completion of all async copies using the **wait_group_events** built-in function before exiting; otherwise the behavior is undefined.

6.12.11 Atomic Functions

The OpenCL C programming language implements the following functions that provide atomic operations on 32-bit signed, unsigned integers and single precision floating-point⁵⁸ to locations in `__global` or `__local` memory.

| Function | Description |
|--|---|
| <pre>int atomic_add (volatile __global int *p, int val) unsigned int atomic_add (volatile __global unsigned int *p, unsigned int val) int atomic_add (volatile __local int *p, int val) unsigned int atomic_add (volatile __local unsigned int *p, unsigned int val)</pre> | <p>Read the 32-bit value (referred to as <i>old</i>) stored at location pointed by <i>p</i>. Compute (<i>old</i> + <i>val</i>) and store result at location pointed by <i>p</i>. The function returns <i>old</i>.</p> |
| <pre>int atomic_sub (volatile __global int *p, int val) unsigned int atomic_sub (volatile __global unsigned int *p, unsigned int val) int atomic_sub (volatile __local int *p, int val) unsigned int atomic_sub (volatile __local unsigned int *p, unsigned int val)</pre> | <p>Read the 32-bit value (referred to as <i>old</i>) stored at location pointed by <i>p</i>. Compute (<i>old</i> - <i>val</i>) and store result at location pointed by <i>p</i>. The function returns <i>old</i>.</p> |
| <pre>int atomic_xchg (volatile __global int *p, int val) unsigned int atomic_xchg (volatile __global unsigned int *p, unsigned int val) float atomic_xchg (volatile __global float *p, float val) int atomic_xchg (volatile __local int *p, int val) unsigned int atomic_xchg (volatile __local unsigned int *p, unsigned int val) float atomic_xchg (volatile __local float *p, float val)</pre> | <p>Swaps the <i>old</i> value stored at location <i>p</i> with new value given by <i>val</i>. Returns <i>old</i> value.</p> |
| <pre>int atomic_inc (volatile __global int *p)</pre> | <p>Read the 32-bit value (referred to as</p> |

⁵⁸ Only the **atomic_xchg** operation is supported for single precision floating-point data type.

| | |
|---|--|
| <pre> unsigned int atomic_inc (volatile __global unsigned int *p) int atomic_inc (volatile __local int *p) unsigned int atomic_inc (volatile __local unsigned int *p) </pre> | <p><i>old</i>) stored at location pointed by <i>p</i>. Compute (<i>old</i> + 1) and store result at location pointed by <i>p</i>. The function returns <i>old</i>.</p> |
| <pre> int atomic_dec (volatile __global int *p) unsigned int atomic_dec (volatile __global unsigned int *p) int atomic_dec (volatile __local int *p) unsigned int atomic_dec (volatile __local unsigned int *p) </pre> | <p>Read the 32-bit value (referred to as <i>old</i>) stored at location pointed by <i>p</i>. Compute (<i>old</i> - 1) and store result at location pointed by <i>p</i>. The function returns <i>old</i>.</p> |
| <pre> int atomic_cmpxchg (volatile __global int *p, int cmp, int val) unsigned int atomic_cmpxchg (volatile __global unsigned int *p, unsigned int cmp, unsigned int val) int atomic_cmpxchg (volatile __local int *p, int cmp, int val) unsigned int atomic_cmpxchg (volatile __local unsigned int *p, unsigned int cmp, unsigned int val) </pre> | <p>Read the 32-bit value (referred to as <i>old</i>) stored at location pointed by <i>p</i>. Compute (<i>old</i> == <i>cmp</i>) ? <i>val</i> : <i>old</i> and store result at location pointed by <i>p</i>. The function returns <i>old</i>.</p> |
| <pre> int atomic_min (volatile __global int *p, int val) unsigned int atomic_min (volatile __global unsigned int *p, unsigned int val) int atomic_min (volatile __local int *p, int val) unsigned int atomic_min (volatile __local unsigned int *p, unsigned int val) </pre> | <p>Read the 32-bit value (referred to as <i>old</i>) stored at location pointed by <i>p</i>. Compute min(<i>old</i>, <i>val</i>) and store minimum value at location pointed by <i>p</i>. The function returns <i>old</i>.</p> |
| <pre> int atomic_max (volatile __global int *p, int val) unsigned int atomic_max (volatile __global unsigned int *p, unsigned int val) int atomic_max (volatile __local int *p, int val) </pre> | <p>Read the 32-bit value (referred to as <i>old</i>) stored at location pointed by <i>p</i>. Compute max(<i>old</i>, <i>val</i>) and store maximum value at location pointed by <i>p</i>. The function returns <i>old</i>.</p> |

| | |
|---|--|
| <pre> unsigned int atomic_max (volatile __local unsigned int *p, unsigned int val) </pre> | |
| <pre> int atomic_and (volatile __global int *p, int val) unsigned int atomic_and (volatile __global unsigned int *p, unsigned int val) int atomic_and (volatile __local int *p, int val) unsigned int atomic_and (volatile __local unsigned int *p, unsigned int val) </pre> | <p>Read the 32-bit value (referred to as <i>old</i>) stored at location pointed by <i>p</i>. Compute (<i>old</i> & val) and store result at location pointed by <i>p</i>. The function returns <i>old</i>.</p> |
| <pre> int atomic_or (volatile __global int *p, int val) unsigned int atomic_or (volatile __global unsigned int *p, unsigned int val) int atomic_or (volatile __local int *p, int val) unsigned int atomic_or (volatile __local unsigned int *p, unsigned int val) </pre> | <p>Read the 32-bit value (referred to as <i>old</i>) stored at location pointed by <i>p</i>. Compute (<i>old</i> val) and store result at location pointed by <i>p</i>. The function returns <i>old</i>.</p> |
| <pre> int atomic_xor (volatile __global int *p, int val) unsigned int atomic_xor (volatile __global unsigned int *p, unsigned int val) int atomic_xor (volatile __local int *p, int val) unsigned int atomic_xor (volatile __local unsigned int *p, unsigned int val) </pre> | <p>Read the 32-bit value (referred to as <i>old</i>) stored at location pointed by <i>p</i>. Compute (<i>old</i> ^ val) and store result at location pointed by <i>p</i>. The function returns <i>old</i>.</p> |

Table 6.19 *Built-in Atomic Functions*

NOTE: The atomic built-in functions that use the **atom_** prefix and are described by the following extensions

- ✚ **cl_khr_global_int32_base_atomics**
- ✚ **cl_khr_global_int32_extended_atomics**
- ✚ **cl_khr_local_int32_base_atomics**
- ✚ **cl_khr_local_int32_extended_atomics**

in sections 9.5 and 9.6 of the OpenCL 1.0 specification are also supported.

6.12.12 Miscellaneous Vector Functions

The OpenCL C programming language implements the following additional built-in vector functions. We use the generic type name *gentypen* (or *gentypem*) to indicate the built-in data types `char{2|4|8|16}`, `uchar{2|4|8|16}`, `short{2|4|8|16}`, `ushort{2|4|8|16}`, `half{2|4|8|16}`⁵⁹, `int{2|4|8|16}`, `uint{2|4|8|16}`, `long{2|4|8|16}`, `ulong{2|4|8|16}`, `float{2|4|8|16}` or `double{2|4|8|16}`⁶⁰ as the type for the arguments unless otherwise stated. We use the generic name *ugentypen* to indicate the built-in unsigned integer data types.

| Function | Description |
|---|---|
| <p><code>int vec_step (gentypen a)</code></p> <p><code>int vec_step (char3 a)</code> <code>int vec_step (uchar3 a)</code> <code>int vec_step (short3 a)</code> <code>int vec_step (ushort3 a)</code> <code>int vec_step (half3 a)</code> <code>int vec_step (int3 a)</code> <code>int vec_step (uint3 a)</code> <code>int vec_step (long3 a)</code> <code>int vec_step (ulong3 a)</code> <code>int vec_step (float3 a)</code> <code>int vec_step (double3 a)</code></p> <p><code>int vec_step (type)</code></p> | <p>The vec_step built-in function takes a built-in scalar or vector data type argument and returns an integer value representing the number of elements in the scalar or vector.</p> <p>For all scalar types, vec_step returns 1.</p> <p>The vec_step built-in functions that take a 3-component vector return 4.</p> <p>vec_step may also take a pure type as an argument, e.g. vec_step(float2)</p> |
| <p><code>gentypen shuffle (gentypem x, ugentypen mask)</code></p> <p><code>gentypen shuffle2 (gentypem x, gentypem y, ugentypen mask)</code></p> | <p>The shuffle and shuffle2 built-in functions construct a permutation of elements from one or two input vectors respectively that are of the same type, returning a vector with the same element type as the input and length that is the same as the shuffle mask. The size of each element in the <i>mask</i> must match the size of each element in the result. For shuffle, only the ilogb(2m-1) least significant bits of each <i>mask</i> element are considered. For shuffle2, only the ilogb(2m-1)+1 least significant bits of each <i>mask</i> element are considered. Other bits in the mask shall be ignored.</p> <p>The elements of the input vectors are numbered from left to right across one or both of the vectors.</p> |

⁵⁹ Only if the **cl_khr_fp16** extension is supported.

⁶⁰ Only if double precision is supported.

| | |
|--|--|
| | <p>For this purpose, the number of elements in a vector is given by <code>vec_step(gentypem)</code>. The shuffle <i>mask</i> operand specifies, for each element of the result vector, which element of the one or two input vectors the result element gets.</p> <p>Examples:</p> <pre>uint4 mask = (uint4)(3, 2, 1, 0); float4 a; float4 r = shuffle(a, mask); // r.s0123 = a.wzyx uint8 mask = (uint8) (0, 1, 2, 3, 4, 5, 6, 7); float4 a, b; float8 r = shuffle2(a, b, mask); // r.s0123 = a.xyzw // r.s4567 = b.xyzw uint4 mask; float8 a; float4 b; b = shuffle(a, mask);</pre> <p>Examples that are not valid are:</p> <pre>uint8 mask; short16 a; short8 b; b = shuffle(a, mask); ← not valid</pre> |
|--|--|

Table 6.20 *Built-in Miscellaneous Vector Functions*

6.12.13 printf

The OpenCL C programming language implements the **printf** function.

| Function | Description |
|--|---|
| int printf (constant char * restrict <i>format</i> , ...) | <p>The printf built-in function writes output to an implementation-defined stream such as stdout under control of the string pointed to by <i>format</i> that specifies how subsequent arguments are converted for output. If there are insufficient arguments for the format, the behavior is undefined. If the format is exhausted while arguments remain, the excess arguments are evaluated (as always) but are otherwise ignored. The printf function returns when the end of the format string is encountered.</p> <p>printf returns 0 if it was executed successfully and -1 otherwise.</p> |

Table 6.21 *Built-in printf Function*

6.12.13.1 printf output synchronization

When the event that is associated with a particular kernel invocation is completed, the output of all printf() calls executed by this kernel invocation is flushed to the implementation-defined output stream. Calling clFinish on a command queue flushes all pending output by printf in previously enqueued and completed commands to the implementation-defined output stream. In the case that printf is executed from multiple work-items concurrently, there is no guarantee of ordering with respect to written data. For example, it is valid for the output of a work-item with a global id (0,0,1) to appear intermixed with the output of a work-item with a global id (0,0,4) and so on.

6.12.13.2 printf format string

The format shall be a character sequence, beginning and ending in its initial shift state. The format is composed of zero or more directives: ordinary characters (not %), which are copied unchanged to the output stream; and conversion specifications, each of which results in fetching zero or more subsequent arguments, converting them, if applicable, according to the corresponding conversion specifier, and then writing the result to the output stream. As format is in the constant address space it must be resolvable at compile time and thus cannot be dynamically created by the executing program, itself.

Each conversion specification is introduced by the character %. After the %, the following appear in sequence:

- + Zero or more *flags* (in any order) that modify the meaning of the conversion specification.
- + An optional minimum *field width*. If the converted value has fewer characters than the field width, it is padded with spaces (by default) on the left (or right, if the left adjustment flag, described later, has been given) to the field width. The field width takes the form of a nonnegative decimal integer.⁶¹⁾
- + An optional *precision* that gives the minimum number of digits to appear for the **d**, **i**, **o**, **u**, **x**, and **X** conversions, the number of digits to appear after the decimal-point character for **a**, **A**, **e**, **E**, **f**, and **F** conversions, the maximum number of significant digits for the **g** and **G** conversions, or the maximum number of bytes to be written for **s** conversions. The precision takes the form of a period (.) followed by an optional decimal integer; if only the period is specified, the precision is taken as zero. If a precision appears with any other conversion specifier, the behavior is undefined.
- + An optional *vector specifier*.
- + A *length modifier* that specifies the size of the argument. The *length modifier* is required with a vector specifier and together specifies the vector type. Implicit conversions between vector types are disallowed (as per *section 6.2.1*). If the *vector specifier* is not specified, the *length modifier* is optional.
- + A *conversion specifier* character that specifies the type of conversion to be applied.

The flag characters and their meanings are:

- The result of the conversion is left-justified within the field. (It is right-justified if this flag is not specified.)
- + The result of a signed conversion always begins with a plus or minus sign. (It begins with a sign only when a negative value is converted if this flag is not specified.)⁶²⁾
- space* If the first character of a signed conversion is not a sign, or if a signed conversion results in no characters, a space is prefixed to the result. If the *space* and + flags both appear, the *space* flag is ignored.
- # The result is converted to an “alternative form”. For **o** conversion, it increases the precision, if and only if necessary, to force the first digit of the result to be a zero (if the value and precision are both 0, a single 0 is printed). For **x** (or **X**) conversion, a nonzero

⁶¹ Note that **0** is taken as a flag, not as the beginning of a field width.

⁶² The results of all floating conversions of a negative zero, and of negative values that round to zero, include a minus sign.

result has *0x* (or *0X*) prefixed to it. For *a*, *A*, *e*, *E*, *f*, *F*, *g*, and *G* conversions, the result of converting a floating-point number always contains a decimal-point character, even if no digits follow it. (Normally, a decimal-point character appears in the result of these conversions only if a digit follows it.) For *g* and *G* conversions, trailing zeros are *not* removed from the result. For other conversions, the behavior is undefined.

- 0** For *d*, *i*, *o*, *u*, *x*, *X*, *a*, *A*, *e*, *E*, *f*, *F*, *g*, and *G* conversions, leading zeros (following any indication of sign or base) are used to pad to the field width rather than performing space padding, except when converting an infinity or NaN. If the **0** and **-** flags both appear, the **0** flag is ignored. For *d*, *i*, *o*, *u*, *x*, and *X* conversions, if a precision is specified, the **0** flag is ignored. For other conversions, the behavior is undefined.

The vector specifier and its meaning is:

- vn** Specifies that a following *a*, *A*, *e*, *E*, *f*, *F*, *g*, *G*, *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to a vector argument, where *n* is the size of the vector and must be 2, 3, 4, 8 or 16.

The vector value is displayed in the following general form:

value1 C value2 C C valuen

where C is a separator character. The value for this separator character is a comma.

If the vector specifier is not used, the length modifiers and their meanings are:

- hh** Specifies that a following *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to a **char** or **uchar** argument (the argument will have been promoted according to the integer promotions, but its value shall be converted to **char** or **uchar** before printing).
- h** Specifies that a following *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to a **short** or **ushort** argument (the argument will have been promoted according to the integer promotions, but its value shall be converted to **short** or **unsigned short** before printing).
- l** (ell) Specifies that a following *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to a **long** or **ulong** argument. The **l** modifier is supported by the full profile. For the embedded profile, the **l** modifier is supported only if 64-bit integers are supported by the device.

If the vector specifier is used, the length modifiers and their meanings are:

- hh** Specifies that a following *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to a **charn** or **ucharn** argument (the argument will not be promoted).
- h** Specifies that a following *d*, *i*, *o*, *u*, *x*, or *X* conversion specifier applies to a **shortn** or **ushortn** argument (the argument will not be promoted); that

a following **a**, **A**, **e**, **E**, **f**, **F**, **g**, or **G** conversion specifier applies to a **halfn**⁶³ argument.

hl This modifier can only be used with the vector specifier. Specifies that a following **d**, **i**, **o**, **u**, **x**, or **X** conversion specifier applies to a **intn** or **uintn** argument; that a following **a**, **A**, **e**, **E**, **f**, **F**, **g**, or **G** conversion specifier applies to a **floatn** argument.

l (ell) Specifies that a following **d**, **i**, **o**, **u**, **x**, or **X** conversion specifier applies to a **longn** or **ulongn** argument; that a following **a**, **A**, **e**, **E**, **f**, **F**, **g**, or **G** conversion specifier applies to a **doublen** argument. The **l** modifier is supported by the full profile. For the embedded profile, the **l** modifier is supported only if 64-bit integers or double-precision floating-point are supported by the device.

If a vector specifier appears without a length modifier, the behavior is undefined. The vector data type described by the vector specifier and length modifier must match the data type of the argument; otherwise the behavior is undefined.

If a length modifier appears with any conversion specifier other than as specified above, the behavior is undefined.

The conversion specifiers and their meanings are:

d,i The **int**, **charn**, **shortn**, **intn** or **longn** argument is converted to signed decimal in the style *[-]dddd*. The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it is expanded with leading zeros. The default precision is 1. The result of converting a zero value with a precision of zero is no characters.

**o,u,
x,X** The **unsigned int**, **ucharn**, **ushortn**, **uintn** or **ulongn** argument is converted to unsigned octal (**o**), unsigned decimal (**u**), or unsigned hexadecimal notation (**x** or **X**) in the style *dddd*; the letters **abcdef** are used for **x** conversion and the letters **ABCDEF** for **X** conversion. The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it is expanded with leading zeros. The default precision is 1. The result of converting a zero value with a precision of zero is no characters.

f,F A **double**, **halfn**, **floatn** or **doublen** argument representing a floating-point number is converted to decimal notation in the style *[-]ddd.ddd*, where the number of digits after the decimal-point character is equal to the precision specification. If the precision is missing, it is taken as 6; if the precision is zero and the **#** flag is not specified, no decimal-point character appears. If a decimal-point character appears, at least one digit appears before it. The value is rounded to the appropriate number of digits. A **double**, **halfn**, **floatn** or **doublen** argument representing an infinity is converted in one of the styles *[-]inf* or *[-]infinity* — which style is implementation-defined. A **double**, **halfn**,

⁶³ Only if the `cl_khr_fp16` extension is supported.

floatn or **doublen** argument representing a NaN is converted in one of the styles *[-/]**nan**(*n-char-sequence*)* — which style, and the meaning of any *n-char-sequence*, is implementation-defined. The **F** conversion specifier produces **INF**, **INFINITY**, or **NAN** instead of **inf**, **infinity**, or **nan**, respectively.⁶⁴⁾

e,E A **double**, **halfn**, **floatn** or **doublen** argument representing a floating-point number is converted in the style *[-/]d.ddd e±dd*, where there is one digit (which is nonzero if the argument is nonzero) before the decimal-point character and the number of digits after it is equal to the precision; if the precision is missing, it is taken as 6; if the precision is zero and the **#** flag is not specified, no decimal-point character appears. The value is rounded to the appropriate number of digits. The **E** conversion specifier produces a number with **E** instead of **e** introducing the exponent. The exponent always contains at least two digits, and only as many more digits as necessary to represent the exponent. If the value is zero, the exponent is zero. A **double**, **halfn**, **floatn** or **doublen** argument representing an infinity or NaN is converted in the style of an **f** or **F** conversion specifier.

g,G A **double**, **halfn**, **floatn** or **doublen** argument representing a floating-point number is converted in style **f** or **e** (or in style **F** or **E** in the case of a **G** conversion specifier), depending on the value converted and the precision. Let *P* equal the precision if nonzero, 6 if the precision is omitted, or 1 if the precision is zero. Then, if a conversion with style **E** would have an exponent of *X*: — if $P > X \geq -4$, the conversion is with style **f** (or **F**) and precision $P - (X + 1)$. — otherwise, the conversion is with style **e** (or **E**) and precision $P - 1$. Finally, unless the **#** flag is used, any trailing zeros are removed from the fractional portion of the result and the decimal-point character is removed if there is no fractional portion remaining. A **double**, **halfn**, **floatn** or **doublen** **e** argument representing an infinity or NaN is converted in the style of an **f** or **F** conversion specifier.

a,A A **double**, **halfn**, **floatn** or **doublen** argument representing a floating-point number is converted in the style *[-/]**0xh.hhhh p±d***, where there is one hexadecimal digit (which is nonzero if the argument is a normalized floating-point number and is otherwise unspecified) before the decimal-point character⁶⁵⁾ and the number of hexadecimal digits after it is equal to the precision; if the precision is missing, then the precision is sufficient for an exact representation of the value; if the precision is zero and the **#** flag is not specified, no decimal point character appears. The letters **abcdef** are used for **a** conversion and the letters **ABCDEF** for **A** conversion. The **A** conversion specifier produces a number with **X** and **P** instead of **x** and **p**. The exponent always contains at least one digit, and only as many more digits as necessary to represent the decimal exponent of 2. If the value is zero, the exponent is zero. A **double**, **halfn**, **floatn** or **doublen** argument representing an infinity or NaN is converted in the style of an **f** or **F** conversion specifier.

⁶⁴ When applied to infinite and NaN values, the **-**, **+**, and *space* flag characters have their usual meaning; the **#** and **0** flag characters have no effect.

⁶⁵ Binary implementations can choose the hexadecimal digit to the left of the decimal-point character so that subsequent digits align to nibble (4-bit) boundaries.

NOTE: The conversion specifiers **e,E,g,G,a,A** convert a float or half argument that is a scalar type to a double only if the double data type is supported. If the double data type is not supported, the argument will be a **float** instead of a **double** and the half type will be converted to a float.

- c** The **int** argument is converted to an **unsigned char**, and the resulting character is written.
- s** The argument shall be a literal string.⁶⁶ Characters from the literal string array are written up to (but not including) the terminating null character. If the precision is specified, no more than that many bytes are written. If the precision is not specified or is greater than the size of the array, the array shall contain a null character.
- p** The argument shall be a pointer to **void**. The pointer can refer to a memory region in the `global`, `constant`, `local` or `private` address space. The value of the pointer is converted to a sequence of printing characters in an implementation-defined manner.
- %** A **%** character is written. No argument is converted. The complete conversion specification shall be **%%**.

If a conversion specification is invalid, the behavior is undefined. If any argument is not the correct type for the corresponding conversion specification, the behavior is undefined.

In no case does a nonexistent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is expanded to contain the conversion result.

For **a** and **A** conversions, the value is correctly rounded to a hexadecimal floating number with the given precision.

A few examples of `printf` are given below:

```
float4 f = (float4) (1.0f, 2.0f, 3.0f, 4.0f);
uchar4 uc = (uchar4) (0xFA, 0xFB, 0xFC, 0xFD);

printf("f4 = %2.2v4hlf\n", f);
printf("uc = %#v4hhx\n", uc);
```

The above two `printf` calls print the following:

```
f4 = 1.00, 2.00, 3.00, 4.00
uc = 0xfa, 0xfb, 0xfc, 0xfd
```

⁶⁶ No special provisions are made for multibyte characters. The behavior of **printf** with the **s** conversion specifier is undefined if the argument value is not a pointer to a literal string.

A few examples of valid use cases of `printf` for the conversion specifier `s` are given below. The argument value must be a pointer to a literal string.

```
kernel void my_kernel( ... )
{
    printf("%s\n", "this is a test string\n");
}
```

A few examples of invalid use cases of `printf` for the conversion specifier `s` are given below:

```
kernel void my_kernel(global char *s, ... )
{
    printf("%s\n", s);

    constant char *p = "this is a test string\n";
    printf("%s\n", p);
    printf("%s\n", &p[3]);
}
```

A few examples of invalid use cases of `printf` where data types given by the vector specifier and length modifier do not match the argument type are given below:

```
kernel void my_kernel(global char *s, ... )
{
    uint2 ui = (uint2)(0x12345678, 0x87654321);
    printf("unsigned short value = (%#v2hx)\n", ui)
    printf("unsigned char value = (%#v2hhx)\n", ui)
}
```

6.12.13.3 Differences between OpenCL C and C99 `printf`

- ✚ The **I** modifier followed by a **c** conversion specifier or **s** conversion specifier is not supported by OpenCL C.
- ✚ The **ll**, **j**, **z**, **t**, and **L** length modifiers are not supported by OpenCL C but are reserved.
- ✚ The **n** conversion specifier is not supported by OpenCL C but is reserved.
- ✚ OpenCL C adds the optional **vn** vector specifier to support printing of vector types.
- ✚ The conversion specifiers **f**, **F**, **e**, **E**, **g**, **G**, **a**, **A** convert a float argument to a double only if the double data type is supported. Refer to the description of `CL_DEVICE_DOUBLE_FP_CONFIG` in *table 4.3*. If the double data type is not supported, the argument will be a float instead of a double.
- ✚ For the embedded profile, the **I** length modifier is supported only if 64-bit integers are

supported.

- ✚ In OpenCL C, **printf** returns 0 if it was executed successfully and -1 otherwise vs. C99 where **printf** returns the number of characters printed or a negative value if an output or encoding error occurred.
- ✚ In OpenCL C, the conversion specifier **s** can only be used for arguments that are literal strings.

6.12.14 Image Read and Write Functions

The built-in functions defined in this section can only be used with image memory objects. An image memory object can be accessed by specific function calls that read from and/or write to specific locations in the image.


Image memory objects that are being read by a kernel should be declared with the `__read_only` qualifier. **write_image** calls to image memory objects declared with the `__read_only` qualifier will generate a compilation error. Image memory objects that are being written to by a kernel should be declared with the `__write_only` qualifier. **read_image** calls to image memory objects declared with the `__write_only` qualifier will generate a compilation error. **read_image** and **write_image** calls to the same image memory object in a kernel are not supported.

The **read_image** calls returns a four component floating-point, integer or unsigned integer color value. The color values returned by **read_image** are identified as `x`, `y`, `z`, `w` where `x` refers to the red component, `y` refers to the green component, `z` refers to the blue component and `w` refers to the alpha component.

6.12.14.1 Samplers

The image read functions take a sampler argument. The sampler can be passed as an argument to the kernel using `clSetKernelArg`, or can be declared in the outermost scope of `kernel` functions, or it can be a constant variable of type `sampler_t` declared in the program source.

Sampler variables in a program are declared to be of type `sampler_t`. A variable of `sampler_t` type declared in the program source must be initialized with a 32-bit unsigned integer constant, which is interpreted as a bit-field specifying the following properties:

-  Addressing Mode
-  Filter Mode
-  Normalized Coordinates

These properties control how elements of an image object are read by **read_image{f|i|ui}**.

Samplers can also be declared as global constants in the program source using the following syntax.

```
const sampler_t          <sampler name> = <value>
                        or
constant sampler_t      <sampler name> = <value>
                        or
__constant sampler_t    <sampler_name> = <value>
```

Note that samplers declared using the `constant` qualifier are not counted towards the maximum number of arguments pointing to the constant address space or the maximum size of the constant address space allowed per device (i.e. `CL_DEVICE_MAX_CONSTANT_ARGS` and `CL_DEVICE_MAX_CONSTANT_BUFFER_SIZE` as described in *table 4.3*).

The sampler fields are described in *table 6.22*.

| Sampler State | Description |
|---------------------|---|
| <normalized coords> | <p>Specifies whether the <i>x</i>, <i>y</i> and <i>z</i> coordinates are passed in as normalized or unnormalized values. This must be a literal value and can be one of the following predefined enums:</p> <p><code>CLK_NORMALIZED_COORDS_TRUE</code> or <code>CLK_NORMALIZED_COORDS_FALSE</code>.</p> <p>The samplers used with an image in multiple calls to <code>read_image{f ui}</code> declared in a kernel must use the same value for <normalized coords>.</p> |
| <addressing mode> | <p>Specifies the image addressing-mode i.e. how out-of-range image coordinates are handled. This must be a literal value and can be one of the following predefined enums:</p> <p><code>CLK_ADDRESS_MIRRORED_REPEAT</code> - Flip the image coordinate at every integer junction. This addressing mode can only be used with normalized coordinates. If normalized coordinates are not used, this addressing mode may generate image coordinates that are undefined.</p> <p><code>CLK_ADDRESS_REPEAT</code> – out-of-range image coordinates are wrapped to the valid range. This addressing mode can only be used with normalized coordinates. If normalized coordinates are not used, this addressing mode may generate image coordinates that are undefined.</p> <p><code>CLK_ADDRESS_CLAMP_TO_EDGE</code> – out-of-range image coordinates are clamped to the extent.</p> <p><code>CLK_ADDRESS_CLAMP</code>⁶⁷ – out-of-range image coordinates will return a border color.</p> |

⁶⁷ This is similar to the `GL_ADDRESS_CLAMP_TO_BORDER` addressing mode.

| | |
|---------------|--|
| | <p>CLK_ADDRESS_NONE – for this addressing mode the programmer guarantees that the image coordinates used to sample elements of the image refer to a location inside the image; otherwise the results are undefined.</p> <p>For 1D and 2D image arrays, the addressing mode applies only to the x and (x, y) coordinates. The addressing mode for the coordinate which specifies the array index is always CLK_ADDRESS_CLAMP_TO_EDGE.</p> |
| <filter mode> | <p>Specifies the filter mode to use. This must be a literal value and can be one of the following predefined enums: CLK_FILTER_NEAREST or CLK_FILTER_LINEAR.</p> <p>Refer to <i>section 8.2</i> for a description of these filter modes.</p> |

Table 6.22 *Sampler Descriptor*

Examples:

```
const sampler_t  samplerA = CLK_NORMALIZED_COORDS_TRUE |
                           CLK_ADDRESS_REPEAT           |
                           CLK_FILTER_NEAREST;
```

samplerA specifies a sampler that uses normalized coordinates, the repeat addressing mode and a nearest filter.

The maximum number of samplers that can be declared in a kernel can be queried using the CL_DEVICE_MAX_SAMPLERS token in **clGetDeviceInfo**.

6.12.14.1.1 Determining the border color

If <addressing mode> in sampler is CLK_ADDRESS_CLAMP, then out-of-range image coordinates return the border color. The border color selected depends on the image channel order and can be one of the following values:

- ✚ If the image channel order is CL_A, CL_INTENSITY, CL_Rx, CL_RA, CL_RGx, CL_RGBx, CL_ARGB, CL_BGRA, or CL_RGBA, the border color is (0.0f, 0.0f, 0.0f, 0.0f).
- ✚ If the image channel order is CL_R, CL_RG, CL_RGB, or CL_LUMINANCE, the border color is (0.0f, 0.0f, 0.0f, 1.0f).

6.12.14.2 Built-in Image Read Functions

The following built-in function calls to read images with a sampler⁶⁸ are supported.

| Function | Description |
|---|---|
| <p>float4 read_imagef (image2d_t <i>image</i>, sampler_t <i>sampler</i>, int2 <i>coord</i>)</p> <p>float4 read_imagef (image2d_t <i>image</i>, sampler_t <i>sampler</i>, float2 <i>coord</i>)</p> | <p>Use the coordinate (<i>coord.x</i>, <i>coord.y</i>) to do an element lookup in the 2D image object specified by <i>image</i>.</p> <p>read_imagef returns floating-point values in the range [0.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to one of the pre-defined packed formats or CL_UNORM_INT8, or CL_UNORM_INT16.</p> <p>read_imagef returns floating-point values in the range [-1.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to CL_SNORM_INT8, or CL_SNORM_INT16.</p> <p>read_imagef returns floating-point values for image objects created with <i>image_channel_data_type</i> set to CL_HALF_FLOAT or CL_FLOAT.</p> <p>The read_imagef calls that take integer coordinates must use a sampler with filter mode set to CLK_FILTER_NEAREST, normalized coordinates set to CLK_NORMALIZED_COORDS_FALSE and addressing mode set to CLK_ADDRESS_CLAMP_TO_EDGE, CLK_ADDRESS_CLAMP or CLK_ADDRESS_NONE; otherwise the values returned are undefined.</p> <p>Values returned by read_imagef for image objects with <i>image_channel_data_type</i> values not specified in the description above are undefined.</p> |
| <p>int4 read_imagei (image2d_t <i>image</i>, sampler_t <i>sampler</i>, int2 <i>coord</i>)</p> <p>int4 read_imagei (image2d_t <i>image</i>,</p> | <p>Use the coordinate (<i>coord.x</i>, <i>coord.y</i>) to do an element lookup in the 2D image object specified by <i>image</i>.</p> <p>read_imagei and read_imageui return</p> |

⁶⁸ The built-in function calls to read images with a sampler are not supported for image1d_buffer_t image types.

| | |
|---|---|
| <p>sampler_t <i>sampler</i>, float2 <i>coord</i>)</p> <p>uint4 read_imageui (image2d_t <i>image</i>, sampler_t <i>sampler</i>, int2 <i>coord</i>)</p> <p>uint4 read_imageui (image2d_t <i>image</i>, sampler_t <i>sampler</i>, float2 <i>coord</i>)</p> | <p>unnormalized signed integer and unsigned integer values respectively. Each channel will be stored in a 32-bit integer.</p> <p>read_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imagei are undefined.</p> <p>read_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imageui are undefined.</p> <p>The read_image{i ui} calls support a nearest filter only. The filter_mode specified in <i>sampler</i> must be set to CLK_FILTER_NEAREST; otherwise the values returned are undefined.</p> <p>Furthermore, the read_image{i ui} calls that take integer coordinates must use a sampler with normalized coordinates set to CLK_NORMALIZED_COORDS_FALSE and addressing mode set to CLK_ADDRESS_CLAMP_TO_EDGE, CLK_ADDRESS_CLAMP or CLK_ADDRESS_NONE; otherwise the values returned are undefined.</p> |
| <p>float4 read_imagef (image3d_t <i>image</i>, sampler_t <i>sampler</i>, int4 <i>coord</i>)</p> <p>float4 read_imagef (image3d_t <i>image</i>, sampler_t <i>sampler</i>, float4 <i>coord</i>)</p> | <p>Use the coordinate (<i>coord.x</i>, <i>coord.y</i>, <i>coord.z</i>) to do an element lookup in the 3D image object specified by <i>image</i>. <i>coord.w</i> is ignored.</p> <p>read_imagef returns floating-point values in the range [0.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to one of the pre-defined packed formats or CL_UNORM_INT8, or</p> |

| | |
|---|--|
| | <p>CL_UNORM_INT16.</p> <p>read_imagef returns floating-point values in the range [-1.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to CL_SNORM_INT8, or CL_SNORM_INT16.</p> <p>read_imagef returns floating-point values for image objects created with <i>image_channel_data_type</i> set to CL_HALF_FLOAT or CL_FLOAT.</p> <p>The read_imagef calls that take integer coordinates must use a sampler with filter mode set to CLK_FILTER_NEAREST, normalized coordinates set to CLK_NORMALIZED_COORDS_FALSE and addressing mode set to CLK_ADDRESS_CLAMP_TO_EDGE, CLK_ADDRESS_CLAMP or CLK_ADDRESS_NONE; otherwise the values returned are undefined.</p> <p>Values returned by read_imagef for image objects with <i>image_channel_data_type</i> values not specified in the description are undefined.</p> |
| <p>int4 read_imagei (image3d_t <i>image</i>, sampler_t <i>sampler</i>, int4 <i>coord</i>)</p> <p>int4 read_imagei (image3d_t <i>image</i>, sampler_t <i>sampler</i>, float4 <i>coord</i>)</p> <p>uint4 read_imageui (image3d_t <i>image</i>, sampler_t <i>sampler</i>, int4 <i>coord</i>)</p> <p>uint4 read_imageui (image3d_t <i>image</i>, sampler_t <i>sampler</i>, float4 <i>coord</i>)</p> | <p>Use the coordinate (<i>coord.x, coord.y, coord.z</i>) to do an element lookup in the 3D image object specified by <i>image</i>. <i>coord.w</i> is ignored.</p> <p>read_imagei and read_imageui return unnormalized signed integer and unsigned integer values respectively. Each channel will be stored in a 32-bit integer.</p> <p>read_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imagei are undefined.</p> <p>read_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values:</p> |

| | |
|---|---|
| | <p>CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imageui are undefined.</p> <p>The read_image{<i>i</i> <i>ui</i>} calls support a nearest filter only. The filter_mode specified in <i>sampler</i> must be set to CLK_FILTER_NEAREST; otherwise the values returned are undefined.</p> <p>Furthermore, the read_image{<i>i</i> <i>ui</i>} calls that take integer coordinates must use a sampler with normalized coordinates set to CLK_NORMALIZED_COORDS_FALSE and addressing mode set to CLK_ADDRESS_CLAMP_TO_EDGE, CLK_ADDRESS_CLAMP or CLK_ADDRESS_NONE; otherwise the values returned are undefined.</p> |
| <p>float4 read_imagef (image2d_array_t <i>image</i>, sampler_t <i>sampler</i>, int4 <i>coord</i>)</p> <p>float4 read_imagef (image2d_array_t <i>image</i>, sampler_t <i>sampler</i>, float4 <i>coord</i>)</p> | <p>Use <i>coord.xy</i> to do an element lookup in the 2D image identified by <i>coord.z</i> in the 2D image array specified by <i>image</i>.</p> <p>read_imagef returns floating-point values in the range [0.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to one of the pre-defined packed formats or CL_UNORM_INT8, or CL_UNORM_INT16.</p> <p>read_imagef returns floating-point values in the range [-1.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to CL_SNORM_INT8, or CL_SNORM_INT16.</p> <p>read_imagef returns floating-point values for image objects created with <i>image_channel_data_type</i> set to CL_HALF_FLOAT or CL_FLOAT.</p> <p>The read_imagef calls that take integer coordinates must use a sampler with filter mode set to CLK_FILTER_NEAREST, normalized coordinates set to CLK_NORMALIZED_COORDS_FALSE and addressing mode set to CLK_ADDRESS_CLAMP_TO_EDGE,</p> |

| | |
|---|--|
| | <p>CLK_ADDRESS_CLAMP or CLK_ADDRESS_NONE; otherwise the values returned are undefined.</p> <p>Values returned by read_imagef for image objects with <i>image_channel_data_type</i> values not specified in the description above are undefined.</p> |
| <p>int4 read_imagei (<i>image2d_array_t image</i>, sampler_t <i>sampler</i>, int4 <i>coord</i>)</p> <p>int4 read_imagei (<i>image2d_array_t image</i>, sampler_t <i>sampler</i>, float4 <i>coord</i>)</p> <p>uint4 read_imageui (image2d_array_t <i>image</i>, sampler_t <i>sampler</i>, int4 <i>coord</i>)</p> <p>uint4 read_imageui (image2d_array_t <i>image</i>, sampler_t <i>sampler</i>, float4 <i>coord</i>)</p> | <p>Use <i>coord.xy</i> to do an element lookup in the 2D image identified by <i>coord.z</i> in the 2D image array specified by <i>image</i>.</p> <p>read_imagei and read_imageui return unnormalized signed integer and unsigned integer values respectively. Each channel will be stored in a 32-bit integer.</p> <p>read_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imagei are undefined.</p> <p>read_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imageui are undefined.</p> <p>The read_image{i ui} calls support a nearest filter only. The filter_mode specified in <i>sampler</i> must be set to CLK_FILTER_NEAREST; otherwise the values returned are undefined.</p> <p>Furthermore, the read_image{i ui} calls that take integer coordinates must use a sampler with normalized coordinates set to CLK_NORMALIZED_COORDS_FALSE and addressing mode set to CLK_ADDRESS_CLAMP_TO_EDGE,</p> |

| | |
|--|--|
| | CLK_ADDRESS_CLAMP or CLK_ADDRESS_NONE; otherwise the values returned are undefined. |
| <p>float4 read_imagef (image1d_t <i>image</i>, sampler_t <i>sampler</i>, int <i>coord</i>)</p> <p>float4 read_imagef (image1d_t <i>image</i>, sampler_t <i>sampler</i>, float <i>coord</i>)</p> | <p>Use <i>coord</i> to do an element lookup in the 1D image object specified by <i>image</i>.</p> <p>read_imagef returns floating-point values in the range [0.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to one of the pre-defined packed formats or CL_UNORM_INT8, or CL_UNORM_INT16.</p> <p>read_imagef returns floating-point values in the range [-1.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to CL_SNORM_INT8, or CL_SNORM_INT16.</p> <p>read_imagef returns floating-point values for image objects created with <i>image_channel_data_type</i> set to CL_HALF_FLOAT or CL_FLOAT.</p> <p>The read_imagef calls that take integer coordinates must use a sampler with filter mode set to CLK_FILTER_NEAREST, normalized coordinates set to CLK_NORMALIZED_COORDS_FALSE and addressing mode set to CLK_ADDRESS_CLAMP_TO_EDGE, CLK_ADDRESS_CLAMP or CLK_ADDRESS_NONE; otherwise the values returned are undefined.</p> <p>Values returned by read_imagef for image objects with <i>image_channel_data_type</i> values not specified in the description above are undefined.</p> |
| <p>int4 read_imagei (image1d_t <i>image</i>, sampler_t <i>sampler</i>, int <i>coord</i>)</p> <p>int4 read_imagei (image1d_t <i>image</i>, sampler_t <i>sampler</i>, float <i>coord</i>)</p> <p>uint4 read_imageui (image1d_t <i>image</i>,</p> | <p>Use <i>coord</i> to do an element lookup in the 1D image object specified by <i>image</i>.</p> <p>read_imagei and read_imageui return unnormalized signed integer and unsigned integer values respectively. Each channel will be stored in a 32-bit integer.</p> <p>read_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8,</p> |

| | |
|---|--|
| <p>sampler_t <i>sampler</i>, int <i>coord</i>)</p> <p>uint4 read_imageui (image1d_t <i>image</i>, sampler_t <i>sampler</i>, float <i>coord</i>)</p> | <p>CL_SIGNED_INT16 and CL_SIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imagei are undefined.</p> <p>read_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imageui are undefined.</p> <p>The read_image{i ui} calls support a nearest filter only. The filter_mode specified in <i>sampler</i> must be set to CLK_FILTER_NEAREST; otherwise the values returned are undefined.</p> <p>Furthermore, the read_image{i ui} calls that take integer coordinates must use a sampler with normalized coordinates set to CLK_NORMALIZED_COORDS_FALSE and addressing mode set to CLK_ADDRESS_CLAMP_TO_EDGE, CLK_ADDRESS_CLAMP or CLK_ADDRESS_NONE; otherwise the values returned are undefined.</p> |
| <p>float4 read_imagef (image1d_array_t <i>image</i>, sampler_t <i>sampler</i>, int2 <i>coord</i>)</p> <p>float4 read_imagef (image1d_array_t <i>image</i>, sampler_t <i>sampler</i>, float2 <i>coord</i>)</p> | <p>Use <i>coord.x</i> to do an element lookup in the 1D image identified by <i>coord.y</i> in the 1D image array specified by <i>image</i>.</p> <p>read_imagef returns floating-point values in the range [0.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to one of the pre- defined packed formats or CL_UNORM_INT8, or CL_UNORM_INT16.</p> <p>read_imagef returns floating-point values in the range [-1.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to CL_SNORM_INT8, or CL_SNORM_INT16.</p> <p>read_imagef returns floating-point values for image</p> |

| | |
|---|---|
| | <p>objects created with <code>image_channel_data_type</code> set to <code>CL_HALF_FLOAT</code> or <code>CL_FLOAT</code>.</p> <p>The read_imagef calls that take integer coordinates must use a sampler with filter mode set to <code>CLK_FILTER_NEAREST</code>, normalized coordinates set to <code>CLK_NORMALIZED_COORDS_FALSE</code> and addressing mode set to <code>CLK_ADDRESS_CLAMP_TO_EDGE</code>, <code>CLK_ADDRESS_CLAMP</code> or <code>CLK_ADDRESS_NONE</code>; otherwise the values returned are undefined.</p> <p>Values returned by read_imagef for image objects with <code>image_channel_data_type</code> values not specified in the description above are undefined.</p> |
| <p><code>int4 read_imagei (image1d_array_t image, sampler_t sampler, int2 coord)</code></p> <p><code>int4 read_imagei (image1d_array_t image, sampler_t sampler, float2 coord)</code></p> <p><code>uint4 read_imageui (image1d_array_t image, sampler_t sampler, int2 coord)</code></p> <p><code>uint4 read_imageui (image1d_array_t image, sampler_t sampler, float2 coord)</code></p> | <p>Use <i>coord.x</i> to do an element lookup in the 1D image identified by <i>coord.y</i> in the 1D image array specified by <i>image</i>.</p> <p>read_imagei and read_imageui return unnormalized signed integer and unsigned integer values respectively. Each channel will be stored in a 32-bit integer.</p> <p>read_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: <code>CL_SIGNED_INT8</code>, <code>CL_SIGNED_INT16</code> and <code>CL_SIGNED_INT32</code>. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imagei are undefined.</p> <p>read_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: <code>CL_UNSIGNED_INT8</code>, <code>CL_UNSIGNED_INT16</code> and <code>CL_UNSIGNED_INT32</code>. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imageui are undefined.</p> <p>The read_image{i ui} calls support a nearest filter only. The filter_mode specified in <i>sampler</i></p> |

| | |
|--|--|
| | <p>must be set to CLK_FILTER_NEAREST; otherwise the values returned are undefined.</p> <p>Furthermore, the read_image{i ui} calls that take integer coordinates must use a sampler with normalized coordinates set to CLK_NORMALIZED_COORDS_FALSE and addressing mode set to CLK_ADDRESS_CLAMP_TO_EDGE, CLK_ADDRESS_CLAMP or CLK_ADDRESS_NONE; otherwise the values returned are undefined.</p> |
| | |

Table 6.23 Built-in Image Read Functions

6.12.14.3 Built-in Image Sampler-less Read Functions

The following built-in function calls to read images with a sampler are supported. The sampler-less read image functions behave exactly as the corresponding read image functions described in *section 6.12.14.2* that take integer coordinates and a sampler with filter mode set to CLK_FILTER_NEAREST, normalized coordinates set to CLK_NORMALIZED_COORDS_FALSE and addressing mode to CLK_ADDRESS_NONE.

| Function | Description |
|---|---|
| float4 read_imagef (image2d_t <i>image</i> , int2 <i>coord</i>) | <p>Use the coordinate (<i>coord.x</i>, <i>coord.y</i>) to do an element lookup in the 2D image object specified by <i>image</i>.</p> <p>read_imagef returns floating-point values in the range [0.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to one of the pre-defined packed formats or CL_UNORM_INT8, or CL_UNORM_INT16.</p> <p>read_imagef returns floating-point values in the range [-1.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to CL_SNORM_INT8, or CL_SNORM_INT16.</p> <p>read_imagef returns floating-point values for image objects created with <i>image_channel_data_type</i> set to CL_HALF_FLOAT or CL_FLOAT.</p> <p>Values returned by read_imagef for image objects</p> |

| | |
|---|--|
| | with <i>image_channel_data_type</i> values not specified in the description above are undefined. |
| <pre>int4 read_imagei (image2d_t image, int2 coord) uint4 read_imageui (image2d_t image, int2 coord)</pre> | <p>Use the coordinate (<i>coord.x</i>, <i>coord.y</i>) to do an element lookup in the 2D image object specified by <i>image</i>.</p> <p>read_imagei and read_imageui return unnormalized signed integer and unsigned integer values respectively. Each channel will be stored in a 32-bit integer.</p> <p>read_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imagei are undefined.</p> <p>read_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imageui are undefined.</p> |
| <pre>float4 read_imagef (image3d_t image, int4 coord)</pre> | <p>Use the coordinate (<i>coord.x</i>, <i>coord.y</i>, <i>coord.z</i>) to do an element lookup in the 3D image object specified by <i>image</i>. <i>coord.w</i> is ignored.</p> <p>read_imagef returns floating-point values in the range [0.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to one of the predefined packed formats or CL_UNORM_INT8, or CL_UNORM_INT16.</p> <p>read_imagef returns floating-point values in the range [-1.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to CL_SNORM_INT8, or CL_SNORM_INT16.</p> |

| | |
|--|---|
| | <p>read_imagef returns floating-point values for image objects created with <i>image_channel_data_type</i> set to CL_HALF_FLOAT or CL_FLOAT.</p> <p>Values returned by read_imagef for image objects with <i>image_channel_data_type</i> values not specified in the description are undefined.</p> |
| <p>int4 read_imagei (image3d_t <i>image</i>, int4 <i>coord</i>)</p> <p>uint4 read_imageui (image3d_t <i>image</i>, int4 <i>coord</i>)</p> | <p>Use the coordinate (<i>coord.x</i>, <i>coord.y</i>, <i>coord.z</i>) to do an element lookup in the 3D image object specified by <i>image</i>. <i>coord.w</i> is ignored.</p> <p>read_imagei and read_imageui return unnormalized signed integer and unsigned integer values respectively. Each channel will be stored in a 32-bit integer.</p> <p>read_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imagei are undefined.</p> <p>read_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imageui are undefined.</p> |
| <p>float4 read_imagef (image2d_array_t <i>image</i>, int4 <i>coord</i>)</p> | <p>Use <i>coord.xy</i> to do an element lookup in the 2D image identified by <i>coord.z</i> in the 2D image array specified by <i>image</i>.</p> <p>read_imagef returns floating-point values in the range [0.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to one of the pre-defined packed formats or CL_UNORM_INT8, or</p> |

| | |
|--|--|
| | <p>CL_UNORM_INT16.</p> <p>read_imagef returns floating-point values in the range [-1.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to CL_SNORM_INT8, or CL_SNORM_INT16.</p> <p>read_imagef returns floating-point values for image objects created with <i>image_channel_data_type</i> set to CL_HALF_FLOAT or CL_FLOAT.</p> <p>Values returned by read_imagef for image objects with <i>image_channel_data_type</i> values not specified in the description above are undefined.</p> |
| <p>int4 read_imagei (image2d_array_t <i>image</i>, int4 <i>coord</i>)</p> <p>uint4 read_imageui (image2d_array_t <i>image</i>, int4 <i>coord</i>)</p> | <p>Use <i>coord.xy</i> to do an element lookup in the 2D image identified by <i>coord.z</i> in the 2D image array specified by <i>image</i>.</p> <p>read_imagei and read_imageui return unnormalized signed integer and unsigned integer values respectively. Each channel will be stored in a 32-bit integer.</p> <p>read_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imagei are undefined.</p> <p>read_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imageui are undefined.</p> |
| <p>float4 read_imagef (image1d_t <i>image</i>, int <i>coord</i>)</p> | <p>Use <i>coord</i> to do an element lookup in the 1D image or 1D image buffer object specified by <i>image</i>.</p> |

| | |
|--|---|
| <pre>float4 read_imagef (image1d_buffer_t image, int coord)</pre> | <p>read_imagef returns floating-point values in the range [0.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to one of the pre-defined packed formats or CL_UNORM_INT8, or CL_UNORM_INT16.</p> <p>read_imagef returns floating-point values in the range [-1.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to CL_SNORM_INT8, or CL_SNORM_INT16.</p> <p>read_imagef returns floating-point values for image objects created with <i>image_channel_data_type</i> set to CL_HALF_FLOAT or CL_FLOAT.</p> <p>Values returned by read_imagef for image objects with <i>image_channel_data_type</i> values not specified in the description above are undefined.</p> |
| <pre>int4 read_imagei (image1d_t image, int coord) uint4 read_imageui (image1d_t image, int coord) int4 read_imagei (image1d_buffer_t image, int coord) uint4 read_imageui (image1d_buffer_t image, int coord)</pre> | <p>Use <i>coord</i> to do an element lookup in the 1D image or 1D image buffer object specified by <i>image</i>.</p> <p>read_imagei and read_imageui return unnormalized signed integer and unsigned integer values respectively. Each channel will be stored in a 32-bit integer.</p> <p>read_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imagei are undefined.</p> <p>read_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imageui</p> |

| | |
|---|---|
| | are undefined. |
| <p>float4 read_imagef (image1d_array_t <i>image</i>, int2 <i>coord</i>)</p> | <p>Use <i>coord.x</i> to do an element lookup in the 1D image identified by <i>coord.y</i> in the 1D image array specified by <i>image</i>.</p> <p>read_imagef returns floating-point values in the range [0.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to one of the pre-defined packed formats or CL_UNORM_INT8, or CL_UNORM_INT16.</p> <p>read_imagef returns floating-point values in the range [-1.0 ... 1.0] for image objects created with <i>image_channel_data_type</i> set to CL_SNORM_INT8, or CL_SNORM_INT16.</p> <p>read_imagef returns floating-point values for image objects created with <i>image_channel_data_type</i> set to CL_HALF_FLOAT or CL_FLOAT.</p> <p>Values returned by read_imagef for image objects with <i>image_channel_data_type</i> values not specified in the description above are undefined.</p> |
| <p>int4 read_imagei (image1d_array_t <i>image</i>, int2 <i>coord</i>)</p> <p>uint4 read_imageui (image1d_array_t <i>image</i>, int2 <i>coord</i>)</p> | <p>Use <i>coord.x</i> to do an element lookup in the 1D image identified by <i>coord.y</i> in the 1D image array specified by <i>image</i>.</p> <p>read_imagei and read_imageui return unnormalized signed integer and unsigned integer values respectively. Each channel will be stored in a 32-bit integer.</p> <p>read_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imagei are undefined.</p> <p>read_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of</p> |

| | |
|--|---|
| | <p>the following values: CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32. If the <i>image_channel_data_type</i> is not one of the above values, the values returned by read_imageui are undefined.</p> |
|--|---|

Table 6.24 Built-in Image Sampler-less Read Functions

6.12.14.4 Built-in Image Write Functions

The following built-in function calls to write images are supported.

| Function | Description |
|--|--|
| <pre>void write_imagef (image2d_t image, int2 coord, float4 color)</pre> | <p>Write <i>color</i> value to location specified by <i>coord.xy</i> in the 2D image object specified by <i>image</i>. Appropriate data format conversion to the specified image format is done before writing the color value. <i>coord.x</i> and <i>coord.y</i> are considered to be unnormalized coordinates and must be in the range 0 ... image width – 1, and 0 ... image height – 1.</p> |
| <pre>void write_imagei (image2d_t image, int2 coord, int4 color)</pre> | <p>write_imagef can only be used with image objects created with <i>image_channel_data_type</i> set to one of the pre-defined packed formats or set to CL_SNORM_INT8, CL_UNORM_INT8, CL_SNORM_INT16, CL_UNORM_INT16, CL_HALF_FLOAT or CL_FLOAT. Appropriate data format conversion will be done to convert channel data from a floating-point value to actual data format in which the channels are stored.</p> |
| <pre>void write_imageui (image2d_t image, int2 coord, uint4 color)</pre> | <p>write_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32.</p> <p>write_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and</p> |

| | |
|---|---|
| | <p>CL_UNSIGNED_INT32.</p> <p>The behavior of write_imagef, write_imagei and write_imageui for image objects created with <i>image_channel_data_type</i> values not specified in the description above or with (x, y) coordinate values that are not in the range (0 ... image width – 1, 0 ... image height – 1), respectively, is undefined.</p> |
| <pre>void write_imagef (image2d_array_t image, int4 coord, float4 color) void write_imagei (image2d_array_t image, int4 coord, int4 color) void write_imageui (image2d_array_t image, int4 coord, uint4 color)</pre> | <p>Write <i>color</i> value to location specified by <i>coord.xy</i> in the 2D image identified by <i>coord.z</i> in the 2D image array specified by <i>image</i>. Appropriate data format conversion to the specified image format is done before writing the color value. <i>coord.x</i>, <i>coord.y</i> and <i>coord.z</i> are considered to be unnormalized coordinates and must be in the range 0 ... image width – 1, 0 ... image height – 1 and 0 ... image number of layers – 1.</p> <p>write_imagef can only be used with image objects created with <i>image_channel_data_type</i> set to one of the pre-defined packed formats or set to CL_SNORM_INT8, CL_UNORM_INT8, CL_SNORM_INT16, CL_UNORM_INT16, CL_HALF_FLOAT or CL_FLOAT. Appropriate data format conversion will be done to convert channel data from a floating-point value to actual data format in which the channels are stored.</p> <p>write_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32.</p> <p>write_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32.</p> <p>The behavior of write_imagef, write_imagei and write_imageui for image objects created with <i>image_channel_data_type</i> values not specified in</p> |

| | |
|---|--|
| | the description above or with (x, y, z) coordinate values that are not in the range $(0 \dots \text{image width} - 1, 0 \dots \text{image height} - 1, 0 \dots \text{image number of layers} - 1)$, respectively, is undefined. |
| <pre>void write_imagef (image1d_t image, int coord, float4 color) void write_imagei (image1d_t image, int coord, int4 color) void write_imageui (image1d_t image, int coord, uint4 color) void write_imagef (image1d_buffer_t image, int coord, float4 color) void write_imagei (image1d_buffer_t image, int coord, int4 color) void write_imageui (image1d_buffer_t image, int coord, uint4 color)</pre> | <p>Write <i>color</i> value to location specified by <i>coord</i> in the 1D image or 1D image buffer object specified by <i>image</i>. Appropriate data format conversion to the specified image format is done before writing the color value. <i>coord</i> is considered to be unnormalized coordinates and must be in the range $0 \dots \text{image width} - 1$.</p> <p>write_imagef can only be used with image objects created with <i>image_channel_data_type</i> set to one of the pre-defined packed formats or set to CL_SNORM_INT8, CL_UNORM_INT8, CL_SNORM_INT16, CL_UNORM_INT16, CL_HALF_FLOAT or CL_FLOAT. Appropriate data format conversion will be done to convert channel data from a floating-point value to actual data format in which the channels are stored.</p> <p>write_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32.</p> <p>write_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32.</p> <p>The behavior of write_imagef, write_imagei and write_imageui for image objects created with <i>image_channel_data_type</i> values not specified in the description above or with coordinate values that is not in the range $(0 \dots \text{image width} - 1)$, is undefined.</p> |
| <pre>void write_imagef (image1d_array_t image,</pre> | Write <i>color</i> value to location specified by <i>coord.x</i> in the 1D image identified by <i>coord.y</i> in the 1D image |

| | |
|--|--|
| <pre> int2 coord, float4 color) void write_imagei (image1d_array_t image, int2 coord, int4 color) void write_imageui (image1d_array_t image, int2 coord, uint4 color) </pre> | <p>array specified by <i>image</i>. Appropriate data format conversion to the specified image format is done before writing the color value. <i>coord.x</i> and <i>coord.y</i> are considered to be unnormalized coordinates and must be in the range 0 ... image width – 1 and 0 ... image number of layers – 1.</p> <p>write_imagef can only be used with image objects created with <i>image_channel_data_type</i> set to one of the pre-defined packed formats or set to CL_SNORM_INT8, CL_UNORM_INT8, CL_SNORM_INT16, CL_UNORM_INT16, CL_HALF_FLOAT or CL_FLOAT. Appropriate data format conversion will be done to convert channel data from a floating-point value to actual data format in which the channels are stored.</p> <p>write_imagei can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32.</p> <p>write_imageui can only be used with image objects created with <i>image_channel_data_type</i> set to one of the following values: CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32.</p> <p>The behavior of write_imagef, write_imagei and write_imageui for image objects created with <i>image_channel_data_type</i> values not specified in the description above or with (<i>x</i>, <i>y</i>) coordinate values that are not in the range (0 ... image width – 1, 0 ... image number of layers – 1), respectively, is undefined.</p> |
|--|--|

Table 6.25 *Built-in Image Write Functions*

6.12.14.5 Built-in Image Query Functions

The following built-in function calls to query image information are supported.

| Function | Description |
|---|--|
| int <code>get_image_width</code> (<code>image1d_t image</code>) int <code>get_image_width</code> (<code>image1d_buffer_t image</code>) int <code>get_image_width</code> (<code>image2d_t image</code>) int <code>get_image_width</code> (<code>image3d_t image</code>) int <code>get_image_width</code> (<code>image1d_array_t image</code>) int <code>get_image_width</code> (<code>image2d_array_t image</code>) | Return the image width in pixels. |
| int <code>get_image_height</code> (<code>image2d_t image</code>) int <code>get_image_height</code> (<code>image3d_t image</code>) int <code>get_image_height</code> (<code>image2d_array_t image</code>) | Return the image height in pixels. |
| int <code>get_image_depth</code> (<code>image3d_t image</code>) | Return the image depth in pixels. |
| int <code>get_image_channel_data_type</code> (<code>image1d_t image</code>) int <code>get_image_channel_data_type</code> (<code>image1d_buffer_t image</code>) int <code>get_image_channel_data_type</code> (<code>image2d_t image</code>) int <code>get_image_channel_data_type</code> (<code>image3d_t image</code>) int <code>get_image_channel_data_type</code> (<code>image1d_array_t image</code>) int <code>get_image_channel_data_type</code> (<code>image2d_array_t image</code>) | Return the channel data type. Valid values are: CLK_SNORM_INT8 CLK_SNORM_INT16 CLK_UNORM_INT8 CLK_UNORM_INT16 CLK_UNORM_SHORT_565 CLK_UNORM_SHORT_555 CLK_UNORM_SHORT_101010 CLK_SIGNED_INT8 CLK_SIGNED_INT16 CLK_SIGNED_INT32 CLK_UNSIGNED_INT8 CLK_UNSIGNED_INT16 CLK_UNSIGNED_INT32 CLK_HALF_FLOAT CLK_FLOAT |
| int <code>get_image_channel_order</code> (<code>image1d_t image</code>) int <code>get_image_channel_order</code> (<code>image1d_buffer_t image</code>) int <code>get_image_channel_order</code> (<code>image2d_t image</code>) int <code>get_image_channel_order</code> (<code>image3d_t image</code>) int <code>get_image_channel_order</code> (<code>image2d_array_t image</code>) | Return the image channel order. Valid values are: CLK_A CLK_R CLK_Rx CLK_RG CLK_RGx CLK_RA CLK_RGB |

| | |
|---|--|
| <code>image1d_array_t image)</code> <code>int get_image_channel_order (</code> <code>image2d_array_t image)</code> | CLK_RGBx CLK_RGBA CLK_ARGB CLK_BGRA CLK_INTENSITY CLK_LUMINANCE |
| <code>int2 get_image_dim (image2d_t image)</code> <code>int2 get_image_dim (</code> <code>image2d_array_t image)</code> | Return the 2D image width and height as an <code>int2</code> type. The width is returned in the <i>x</i> component, and the height in the <i>y</i> component. |
| <code>int4 get_image_dim (image3d_t image)</code> | Return the 3D image width, height, and depth as an <code>int4</code> type. The width is returned in the <i>x</i> component, height in the <i>y</i> component, depth in the <i>z</i> component and the <i>w</i> component is 0. |
| <code>size_t get_image_array_size(</code> <code>image2d_array_t image)</code> | Return the number of images in the 2D image array. |
| <code>size_t get_image_array_size(</code> <code>image1d_array_t image)</code> | Return the number of images in the 1D image array. |

Table 6.26 Built-in Image Query Functions

The values returned by `get_image_channel_data_type` and `get_image_channel_order` as specified in *table 6.26* with the CLK_ prefixes correspond to the CL_ prefixes used to describe the image channel order and data type in *tables 5.4* and *5.5*. For example, both CL_UNORM_INT8 and CLK_UNORM_INT8 refer to an image channel data type that is an unnormalized unsigned 8-bit integer.

6.12.14.6 Mapping image channels to color values returned by `read_image` and color values passed to `write_image` to image channels

The following table describes the mapping of the number of channels of an image element to the appropriate components in the `float4`, `int4` or `uint4` vector data type for the color values returned by `read_image{f|i|ui}` or supplied to `write_image{f|i|ui}`. The unmapped components will be set to 0.0 for red, green and blue channels and will be set to 1.0 for the alpha channel.

| Channel Order | float4, int4 or uint4 components of channel data |
|---------------------------|--|
| CL_R, CL_Rx | (r, 0.0, 0.0, 1.0) |
| CL_A | (0.0, 0.0, 0.0, a) |
| CL_RG, CL_RGx | (r, g, 0.0, 1.0) |
| CL_RA | (r, 0.0, 0.0, a) |
| CL_RGB, CL_RGBx | (r, g, b, 1.0) |
| CL_RGBA, CL_BGRA, CL_ARGB | (r, g, b, a) |

| | |
|--------------|----------------|
| CL_INTENSITY | (I, I, I, I) |
| CL_LUMINANCE | (L, L, L, 1.0) |

NOTE: A kernel that uses a sampler with the CL_ADDRESS_CLAMP addressing mode with multiple images may result in additional samplers being used internally by an implementation. If the same sampler is used with multiple images called via **read_image{f|i|ui}**, then it is possible that an implementation may need to allocate an additional sampler to handle the different border color values that may be needed depending on the image formats being used. These implementation allocated samplers will count against the maximum sampler values supported by the device and given by CL_DEVICE_MAX_SAMPLERS. Enqueuing a kernel that requires more samplers than the implementation can support will result in a CL_OUT_OF_RESOURCES error being returned.





7. OpenCL Numerical Compliance

This section describes features of the C99 and IEEE 754 standards that must be supported by all OpenCL compliant devices.

This section describes the functionality that must be supported by all OpenCL devices for single precision floating-point numbers. Currently, only single precision floating-point is a requirement. Double precision floating-point is an optional feature.

7.1 Rounding Modes

Floating-point calculations may be carried out internally with extra precision and then rounded to fit into the destination type. IEEE 754 defines four possible rounding modes:

-  Round to nearest even
-  Round toward $+\infty$
-  Round toward $-\infty$
-  Round toward zero

Round to nearest even is currently the only rounding mode required by the OpenCL specification for single precision and double precision operations and is therefore the default rounding mode. In addition, only static selection of rounding mode is supported. Dynamically reconfiguring the rounding modes as specified by the IEEE 754 spec is unsupported.

7.2 INF, NaN and Denormalized Numbers

INF and NaNs must be supported. Support for signaling NaNs is not required.

Support for denormalized numbers with single precision floating-point is optional. Denormalized single precision floating-point numbers passed as input or produced as the output of single precision floating-point operations such as add, sub, mul, divide, and the functions defined in *sections 6.11.2* (math functions), *6.11.4* (common functions) and *6.11.5* (geometric functions) may be flushed to zero.

7.3 Floating-Point Exceptions

Floating-point exceptions are disabled in OpenCL. The result of a floating-point exception must match the IEEE 754 spec for the exceptions not enabled case. Whether and when the implementation sets floating-point flags or raises floating-point exceptions is implementation-defined. This standard provides no method for querying, clearing or setting floating-point flags or trapping raised exceptions. Due to non-performance, non-portability of trap mechanisms and the impracticality of servicing precise exceptions in a vector context (especially on heterogeneous hardware), such features are discouraged.

Implementations that nevertheless support such operations through an extension to the standard shall initialize with all exception flags cleared and the exception masks set so that exceptions raised by arithmetic operations do not trigger a trap to be taken. If the underlying work is reused by the implementation, the implementation is however not responsible for reclearing the flags or resetting exception masks to default values before entering the kernel. That is to say that kernels that do not inspect flags or enable traps are licensed to expect that their arithmetic will not trigger a trap. Those kernels that do examine flags or enable traps are responsible for clearing flag state and disabling all traps before returning control to the implementation. Whether or when the underlying work-item (and accompanying global floating-point state if any) is reused is implementation-defined.

The expressions **math_errorhandling** and **MATH_ERREXCEPT** are reserved for use by this standard, but not defined. Implementations that extend this specification with support for floating-point exceptions shall define **math_errorhandling** and **MATH_ERREXCEPT** per ISO / IEC 9899 : TC2.

7.4 Relative Error as ULPs

In this section we discuss the maximum relative error defined as `ulp` (units in the last place). Addition, subtraction, multiplication, fused multiply-add and conversion between integer and a single precision floating-point format are IEEE 754 compliant and are therefore correctly rounded. Conversion between floating-point formats and explicit conversions specified in *section 6.2.3* must be correctly rounded.

The ULP is defined as follows:

```
If x is a real number that lies between two finite
consecutive floating-point numbers a and b, without being
equal to one of them, then ulp(x) = |b - a|, otherwise
ulp(x) is the distance between the two non-equal finite
floating-point numbers nearest x. Moreover, ulp(NaN) is
NaN.
```

Attribution: This definition was taken with consent from Jean-Michel Muller with slight clarification for behavior at zero. Refer to <ftp://ftp.inria.fr/INRIA/publication/publi-pdf/RR/RR-5504.pdf>.

Table 7.1⁶⁹ describes the minimum accuracy of single precision floating-point arithmetic operations given as ULP values. The reference value used to compute the ULP value of an arithmetic operation is the infinitely precise result.

| Function | Min Accuracy - ULP values ⁷⁰ |
|-----------------|---|
| $x + y$ | Correctly rounded |
| $x - y$ | Correctly rounded |
| $x * y$ | Correctly rounded |
| $1.0 / x$ | ≤ 2.5 ulp |
| x / y | ≤ 2.5 ulp |
| | |
| acos | ≤ 4 ulp |
| acospi | ≤ 5 ulp |
| asin | ≤ 4 ulp |
| asinpi | ≤ 5 ulp |
| atan | ≤ 5 ulp |
| atan2 | ≤ 6 ulp |
| atanpi | ≤ 5 ulp |
| atan2pi | ≤ 6 ulp |
| acosh | ≤ 4 ulp |
| asinh | ≤ 4 ulp |
| atanh | ≤ 5 ulp |
| cbrt | ≤ 2 ulp |
| ceil | Correctly rounded |
| copysign | 0 ulp |
| cos | ≤ 4 ulp |
| cosh | ≤ 4 ulp |
| cospi | ≤ 4 ulp |
| erfc | ≤ 16 ulp |
| erf | ≤ 16 ulp |
| exp | ≤ 3 ulp |
| exp2 | ≤ 3 ulp |
| exp10 | ≤ 3 ulp |
| expm1 | ≤ 3 ulp |
| fabs | 0 ulp |
| fdim | Correctly rounded |
| floor | Correctly rounded |

⁶⁹ The ULP values for built-in math functions **lgamma** and **lgamma_r** is currently undefined.

⁷⁰ 0 ulp is used for math functions that do not require rounding.

| | |
|--------------------|-------------------------------------|
| fma | Correctly rounded |
| fmax | 0 ulp |
| fmin | 0 ulp |
| fmod | 0 ulp |
| fract | Correctly rounded |
| frexp | 0 ulp |
| hypot | <= 4 ulp |
| ilogb | 0 ulp |
| ldexp | Correctly rounded |
| log | <= 3 ulp |
| log2 | <= 3 ulp |
| log10 | <= 3 ulp |
| log1p | <= 2 ulp |
| logb | 0 ulp |
| mad | Any value allowed (infinite ulp) |
| maxmag | 0 ulp |
| minmag | 0 ulp |
| modf | 0 ulp |
| nan | 0 ulp |
| nextafter | 0 ulp |
| pow(x, y) | <= 16 ulp |
| pown(x, y) | <= 16 ulp |
| powr(x, y) | <= 16 ulp |
| remainder | 0 ulp |
| remquo | 0 ulp |
| rint | Correctly rounded |
| rootn | <= 16 ulp |
| round | Correctly rounded |
| rsqrt | <= 2 ulp |
| sin | <= 4 ulp |
| sincos | <= 4 ulp for sine and cosine values |
| sinh | <= 4 ulp |
| sinpi | <= 4 ulp |
| sqrt | <= 3 ulp |
| tan | <= 5 ulp |
| tanh | <= 5 ulp |
| tanpi | <= 6 ulp |
| tgamma | <= 16 ulp |
| trunc | Correctly rounded |
| | |
| half_cos | <= 8192 ulp |
| half_divide | <= 8192 ulp |
| half_exp | <= 8192 ulp |
| half_exp2 | <= 8192 ulp |
| half_exp10 | <= 8192 ulp |

| | |
|----------------------|------------------------|
| half_log | <= 8192 ulp |
| half_log2 | <= 8192 ulp |
| half_log10 | <= 8192 ulp |
| half_powr | <= 8192 ulp |
| half_recip | <= 8192 ulp |
| half_rsqr | <= 8192 ulp |
| half_sin | <= 8192 ulp |
| half_sqrt | <= 8192 ulp |
| half_tan | <= 8192 ulp |
| | |
| native_cos | Implementation-defined |
| native_divide | Implementation-defined |
| native_exp | Implementation-defined |
| native_exp2 | Implementation-defined |
| native_exp10 | Implementation-defined |
| native_log | Implementation-defined |
| native_log2 | Implementation-defined |
| native_log10 | Implementation-defined |
| native_powr | Implementation-defined |
| native_recip | Implementation-defined |
| native_rsqr | Implementation-defined |
| native_sin | Implementation-defined |
| native_sqrt | Implementation-defined |
| native_tan | Implementation-defined |

Table 7.1 *ULP values for single precision built-in math functions*

Table 7.2 describes the minimum accuracy of double precision floating-point arithmetic operations given as ULP values. The reference value used to compute the ULP value of an arithmetic operation is the infinitely precise result.

| Function | Min Accuracy - ULP values⁷¹ |
|-----------------|---|
| $x + y$ | Correctly rounded |
| $x - y$ | Correctly rounded |
| $x * y$ | Correctly rounded |
| $1.0 / x$ | Correctly rounded |
| x / y | Correctly rounded |
| | |
| acos | <= 4 ulp |
| acospi | <= 5 ulp |
| asin | <= 4 ulp |
| asinpi | <= 5 ulp |
| atan | <= 5 ulp |

⁷¹ 0 ulp is used for math functions that do not require rounding.

| | |
|-------------------|----------------------------------|
| atan2 | <= 6 ulp |
| atanpi | <= 5 ulp |
| atan2pi | <= 6 ulp |
| acosh | <= 4 ulp |
| asinh | <= 4 ulp |
| atanh | <= 5 ulp |
| cbrt | <= 2 ulp |
| ceil | Correctly rounded |
| copysign | 0 ulp |
| cos | <= 4 ulp |
| cosh | <= 4 ulp |
| cospi | <= 4 ulp |
| erfc | <= 16 ulp |
| erf | <= 16 ulp |
| exp | <= 3 ulp |
| exp2 | <= 3 ulp |
| exp10 | <= 3 ulp |
| expm1 | <= 3 ulp |
| fabs | 0 ulp |
| fdim | Correctly rounded |
| floor | Correctly rounded |
| fma | Correctly rounded |
| fmax | 0 ulp |
| fmin | 0 ulp |
| fmod | 0 ulp |
| fract | Correctly rounded |
| frexp | 0 ulp |
| hypot | <= 4 ulp |
| ilogb | 0 ulp |
| ldexp | Correctly rounded |
| log | <= 3 ulp |
| log2 | <= 3 ulp |
| log10 | <= 3 ulp |
| log1p | <= 2 ulp |
| logb | 0 ulp |
| mad | Any value allowed (infinite ulp) |
| maxmag | 0 ulp |
| minmag | 0 ulp |
| modf | 0 ulp |
| nan | 0 ulp |
| nextafter | 0 ulp |
| pow(x, y) | <= 16 ulp |
| pown(x, y) | <= 16 ulp |
| powr(x, y) | <= 16 ulp |
| remainder | 0 ulp |

| | |
|---------------|---|
| remquo | 0 ulp |
| rint | Correctly rounded |
| rootn | ≤ 16 ulp |
| round | Correctly rounded |
| rsqrt | ≤ 2 ulp |
| sin | ≤ 4 ulp |
| sincos | ≤ 4 ulp for sine and cosine values |
| sinh | ≤ 4 ulp |
| sinpi | ≤ 4 ulp |
| sqrt | Correctly rounded |
| tan | ≤ 5 ulp |
| tanh | ≤ 5 ulp |
| tanpi | ≤ 6 ulp |
| tgamma | ≤ 16 ulp |
| trunc | Correctly rounded |

Table 7.2 *ULP values for double precision built-in math functions*

7.5 Edge Case Behavior

The edge case behavior of the math functions (*section 6.12.2*) shall conform to sections F.9 and G.6 of ISO/IEC 9899:TC 2 (commonly known as C99, TC2), except where noted below in *section 7.5.1*.

7.5.1 Additional Requirements Beyond C99 TC2

Functions that return a NaN with more than one NaN operand shall return one of the NaN operands. Functions that return a NaN operand may silence the NaN if it is a signaling NaN. A non-signaling NaN shall be converted to a non-signaling NaN. A signaling NaN shall be converted to a NaN, and should be converted to a non-signaling NaN. How the rest of the NaN payload bits or the sign of NaN is converted is undefined.

half_<funcname> functions behave identically to the function of the same name without the **half_** prefix. They must conform to the same edge case requirements (see sections F.9 and G.6 of C99, TC2). For other cases, except where otherwise noted, these single precision functions are permitted to have up to 8192 ulps of error (as measured in the single precision result), although better accuracy is encouraged.

The usual allowances for rounding error (*section 7.4*) or flushing behavior (*section 7.5.3*) shall not apply for those values for which *section F.9* of C99, TC2, or *sections 7.5.1* and *7.5.3* below (and similar sections for other floating-point precisions) prescribe a result (e.g. **ceil** ($-1 < x < 0$) returns -0). Those values shall produce exactly the prescribed answers, and no

other. Where the \pm symbol is used, the sign shall be preserved. For example, $\sin(\pm 0) = \pm 0$ shall be interpreted to mean $\sin(+0)$ is $+0$ and $\sin(-0)$ is -0 .

acospi (1) = $+0$.

acospi (x) returns a NaN for $|x| > 1$.

asinpi (± 0) = ± 0 .

asinpi (x) returns a NaN for $|x| > 1$.

atanpi (± 0) = ± 0 .

atanpi ($\pm\infty$) = ± 0.5 .

atan2pi (± 0 , -0) = ± 1 .

atan2pi (± 0 , $+0$) = ± 0 .

atan2pi (± 0 , x) returns ± 1 for $x < 0$.

atan2pi (± 0 , x) returns ± 0 for $x > 0$.

atan2pi (y , ± 0) returns -0.5 for $y < 0$.

atan2pi (y , ± 0) returns 0.5 for $y > 0$.

atan2pi ($\pm y$, $-\infty$) returns ± 1 for finite $y > 0$.

atan2pi ($\pm y$, $+\infty$) returns ± 0 for finite $y > 0$.

atan2pi ($\pm\infty$, x) returns ± 0.5 for finite x .

atan2pi ($\pm\infty$, $-\infty$) returns ± 0.75 .

atan2pi ($\pm\infty$, $+\infty$) returns ± 0.25 .

ceil ($-1 < x < 0$) returns -0 .

cospi (± 0) returns 1

cospi ($n + 0.5$) is $+0$ for any integer n where $n + 0.5$ is representable.

cospi ($\pm\infty$) returns a NaN.

exp10 (± 0) returns 1.

exp10 ($-\infty$) returns $+0$.

exp10 ($+\infty$) returns $+\infty$.

distance (x , y) calculates the distance from x to y without overflow or extraordinary precision loss due to underflow.

fdim (any, NaN) returns NaN.

fdim (NaN, any) returns NaN.

fmod (± 0 , NaN) returns NaN.

frexp ($\pm\infty$, exp) returns $\pm\infty$ and stores 0 in exp .

frexp (NaN, exp) returns the NaN and stores 0 in exp .

fract (x , $iptr$) shall not return a value greater than or equal to 1.0, and shall not return a value less than 0.

fract (+0, $iptr$) returns +0 and +0 in $iptr$.

fract (-0, $iptr$) returns -0 and -0 in $iptr$.

fract (+inf, $iptr$) returns +0 and +inf in $iptr$.

fract (-inf, $iptr$) returns -0 and -inf in $iptr$.

fract (NaN, $iptr$) returns the NaN and NaN in $iptr$.

length calculates the length of a vector without overflow or extraordinary precision loss due to underflow.

lgamma_r (x , $signp$) returns 0 in $signp$ if x is zero or a negative integer.

nextafter (-0, $y > 0$) returns smallest positive denormal value.

nextafter (+0, $y < 0$) returns smallest negative denormal value.

normalize shall reduce the vector to unit length, pointing in the same direction without overflow or extraordinary precision loss due to underflow.

normalize (v) returns v if all elements of v are zero.

normalize (v) returns a vector full of NaNs if any element is a NaN.

normalize (v) for which any element in v is infinite shall proceed as if the elements in v were replaced as follows:

```
for( i = 0; i < sizeof(v) / sizeof(v[0]); i++)  
    v[i] = isinf(v[i]) ? copysign(1.0, v[i]) : 0.0 * v [i];
```

pow (± 0 , $-\infty$) returns $+\infty$

pown (x , 0) is 1 for any x , even zero, NaN or infinity.

pown (± 0 , n) is $\pm\infty$ for odd $n < 0$.

pown (± 0 , n) is $+\infty$ for even $n < 0$.

pown (± 0 , n) is +0 for even $n > 0$.

pown (± 0 , n) is ± 0 for odd $n > 0$.

powr (x , ± 0) is 1 for finite $x > 0$.

powr (± 0 , y) is $+\infty$ for finite $y < 0$.

powr (± 0 , $-\infty$) is $+\infty$.

powr (± 0 , y) is +0 for $y > 0$.

powr (+1, y) is 1 for finite y .

powr (x , y) returns NaN for $x < 0$.

powr (± 0 , ± 0) returns NaN.

powr ($+\infty$, ± 0) returns NaN.

powr (+1, $\pm\infty$) returns NaN.

powr (x , NaN) returns the NaN for $x \geq 0$.

powr (NaN, y) returns the NaN.

rint ($-0.5 \leq x < 0$) returns -0 .

remquo ($x, y, \&quo$) returns a NaN and 0 in *quo* if x is $\pm\infty$, or if y is 0 and the other argument is non-NaN or if either argument is a NaN.

rootn ($\pm 0, n$) is $\pm\infty$ for odd $n < 0$.

rootn ($\pm 0, n$) is $+\infty$ for even $n < 0$.

rootn ($\pm 0, n$) is $+0$ for even $n > 0$.

rootn ($\pm 0, n$) is ± 0 for odd $n > 0$.

rootn (x, n) returns a NaN for $x < 0$ and n is even.

rootn ($x, 0$) returns a NaN.

round ($-0.5 < x < 0$) returns -0 .

sinpi (± 0) returns ± 0 .

sinpi ($+n$) returns $+0$ for positive integers n .

sinpi ($-n$) returns -0 for negative integers n .

sinpi ($\pm\infty$) returns a NaN.

tanpi (± 0) returns ± 0 .

tanpi ($\pm\infty$) returns a NaN.

tanpi (n) is **copysign**($0.0, n$) for even integers n .

tanpi (n) is **copysign**($0.0, -n$) for odd integers n .

tanpi ($n + 0.5$) for even integer n is $+\infty$ where $n + 0.5$ is representable.

tanpi ($n + 0.5$) for odd integer n is $-\infty$ where $n + 0.5$ is representable.

trunc ($-1 < x < 0$) returns -0 .

7.5.2 Changes to C99 TC2 Behavior

modf behaves as though implemented by:

```
gentype modf ( gentype value, gentype *iptr )
{
    *iptr = trunc( value );
    return copysign( isinf( value ) ? 0.0 : value - *iptr, value );
}
```

rint always rounds according to round to nearest even rounding mode even if the caller is in some other rounding mode.

7.5.3 Edge Case Behavior in Flush To Zero Mode

If denormals are flushed to zero, then a function may return one of four results:

1. Any conforming result for non-flush-to-zero mode
2. If the result given by 1. is a sub-normal before rounding, it may be flushed to zero
3. Any non-flushed conforming result for the function if one or more of its sub-normal operands are flushed to zero.
4. If the result of 3. is a sub-normal before rounding, the result may be flushed to zero.

In each of the above cases, if an operand or result is flushed to zero, the sign of the zero is undefined.

If subnormals are flushed to zero, a device may choose to conform to the following edge cases for **nextafter** instead of those listed in *section 7.5.1*:

nextafter (+smallest normal, $y < +\text{smallest normal}$) = +0.

nextafter (-smallest normal, $y > -\text{smallest normal}$) = -0.

nextafter (-0, $y > 0$) returns smallest positive normal value.

nextafter (+0, $y < 0$) returns smallest negative normal value.

For clarity, subnormals or denormals are defined to be the set of representable numbers in the range $0 < x < \text{TYPE_MIN}$ and $-\text{TYPE_MIN} < x < -0$. They do not include ± 0 . A non-zero number is said to be sub-normal before rounding if after normalization, its radix-2 exponent is less than $(\text{TYPE_MIN_EXP} - 1)$.⁷²

⁷² Here `TYPE_MIN` and `TYPE_MIN_EXP` should be substituted by constants appropriate to the floating-point type under consideration, such as `FLT_MIN` and `FLT_MIN_EXP` for float.

8. Image Addressing and Filtering

Let w_t , h_t and d_t be the width, height (or image array size for a 1D image array) and depth (or image array size for a 2D image array) of the image in pixels. Let `coord.xy` also referred to as (s, t) or `coord.xyz` also referred to as (s, t, r) be the coordinates specified to `read_image{f|i|ui}`. The sampler specified in `read_image{f|i|ui}` is used to determine how to sample the image and return an appropriate color.

8.1 Image Coordinates

This affects the interpretation of image coordinates. If image coordinates specified to `read_image{f|i|ui}` are normalized (as specified in the sampler), the s , t , and r coordinate values are multiplied by w_t , h_t , and d_t respectively to generate the unnormalized coordinate values.

Let (u, v, w) represent the unnormalized image coordinate values.

8.2 Addressing and Filter Modes

We first describe how the addressing and filter modes are applied to generate the appropriate sample locations to read from the image if the addressing mode is not `CLK_ADDRESS_REPEAT` nor `CLK_ADDRESS_MIRRORED_REPEAT`.

After generating the image coordinate (u, v, w) we apply the appropriate addressing and filter mode to generate the appropriate sample locations to read from the image.

If values in (u, v, w) are INF or NaN, the behavior of `read_image{f|i|ui}` is undefined.

Filter Mode = `CLK_FILTER_NEAREST`

When filter mode is `CLK_FILTER_NEAREST`, the image element in the image that is nearest (in Manhattan distance) to that specified by (u, v, w) is obtained. This means the image element at location (i, j, k) becomes the image element value, where

```
i = address_mode((int) floor(u))
j = address_mode((int) floor(v))
k = address_mode((int) floor(w))
```

For a 3D image, the image element at location (i, j, k) becomes the color value. For a 2D image, the image element at location (i, j) becomes the color value.

Table 8.1 describes the `address_mode` function.

| Addressing Mode | Result of <code>address_mode(coord)</code> |
|--|--|
| <code>CLK_ADDRESS_CLAMP_TO_EDGE</code> | <code>clamp (coord, 0, size - 1)</code> |
| <code>CLK_ADDRESS_CLAMP</code> | <code>clamp (coord, -1, size)</code> |
| <code>CLK_ADDRESS_NONE</code> | <code>coord</code> |

Table 8.1 *Addressing modes to generate texel location.*

The `size` term in table 8.1 is w_t for u , h_t for v and d_t for w .

The `clamp` function used in *table 8.1* is defined as:

$$\text{clamp}(a, b, c) = \text{return } (a < b) ? b : ((a > c) ? c : a)$$

If the selected texel location (i, j, k) refers to a location outside the image, the border color is used as the color value for this texel.

Filter Mode = `CLK_FILTER_LINEAR`

When filter mode is `CLK_FILTER_LINEAR`, a 2×2 square of image elements for a 2D image or a $2 \times 2 \times 2$ cube of image elements for a 3D image is selected. This 2×2 square or $2 \times 2 \times 2$ cube is obtained as follows.

Let

$$\begin{aligned} i_0 &= \text{address_mode}((\text{int})\text{floor}(u - 0.5)) \\ j_0 &= \text{address_mode}((\text{int})\text{floor}(v - 0.5)) \\ k_0 &= \text{address_mode}((\text{int})\text{floor}(w - 0.5)) \\ i_1 &= \text{address_mode}((\text{int})\text{floor}(u - 0.5) + 1) \\ j_1 &= \text{address_mode}((\text{int})\text{floor}(v - 0.5) + 1) \\ k_1 &= \text{address_mode}((\text{int})\text{floor}(w - 0.5) + 1) \\ a &= \text{frac}(u - 0.5) \\ b &= \text{frac}(v - 0.5) \\ c &= \text{frac}(w - 0.5) \end{aligned}$$

where `frac(x)` denotes the fractional part of x and is computed as $x - \text{floor}(x)$.

For a 3D image, the image element value is found as

$$\begin{aligned} T &= (1 - a) * (1 - b) * (1 - c) * T_{i_0j_0k_0} \\ &\quad + a * (1 - b) * (1 - c) * T_{i_1j_0k_0} \\ &\quad + (1 - a) * b * (1 - c) * T_{i_0j_1k_0} \\ &\quad + a * b * (1 - c) * T_{i_1j_1k_0} \end{aligned}$$

$$\begin{aligned}
& + (1 - a) * (1 - b) * c * T_{i0j0k1} \\
& + a * (1 - b) * c * T_{i1j0k1} \\
& + (1 - a) * b * c * T_{i0j1k1} \\
& + a * b * c * T_{i1j1k1}
\end{aligned}$$

where T_{ijk} is the image element at location (i, j, k) in the 3D image.

For a 2D image, the image element value is found as

$$\begin{aligned}
T & = (1 - a) * (1 - b) * T_{i0j0} \\
& + a * (1 - b) * T_{i1j0} \\
& + (1 - a) * b * T_{i0j1} \\
& + a * b * T_{i1j1}
\end{aligned}$$

where T_{ij} is the image element at location (i, j) in the 2D image.

If any of the selected T_{ijk} or T_{ij} in the above equations refers to a location outside the image, the border color is used as the color value for T_{ijk} or T_{ij} .

We now discuss how the addressing and filter modes are applied to generate the appropriate sample locations to read from the image if the addressing mode is CLK_ADDRESS_REPEAT.

If values in (s, t, r) are INF or NaN, the behavior of the built-in image read functions is undefined.

Filter Mode = CLK_FILTER_NEAREST

When filter mode is CLK_FILTER_NEAREST, the image element at location (i, j, k) becomes the image element value, with i, j and k computed as

$$\begin{aligned}
u & = (s - \text{floor}(s)) * w_t \\
i & = (\text{int})\text{floor}(u) \\
\text{if } (i > w_t - 1) \\
& \quad i = i - w_t \\
\\
v & = (t - \text{floor}(t)) * h_t \\
j & = (\text{int})\text{floor}(v) \\
\text{if } (j > h_t - 1) \\
& \quad j = j - h_t \\
\\
w & = (r - \text{floor}(r)) * d_t \\
k & = (\text{int})\text{floor}(w) \\
\text{if } (k > d_t - 1)
\end{aligned}$$

$$k = k - d_t$$

For a 3D image, the image element at location (i, j, k) becomes the color value. For a 2D image, the image element at location (i, j) becomes the color value.

Filter Mode = CLK_FILTER_LINEAR

When filter mode is CLK_FILTER_LINEAR, a 2×2 square of image elements for a 2D image or a $2 \times 2 \times 2$ cube of image elements for a 3D image is selected. This 2×2 square or $2 \times 2 \times 2$ cube is obtained as follows.

Let

```

u = (s - floor(s)) * w_t
i0 = (int)floor(u - 0.5)
i1 = i0 + 1
if (i0 < 0)
    i0 = w_t + i0
if (i1 > w_t - 1)
    i1 = i1 - w_t

```

```

v = (t - floor(t)) * h_t
j0 = (int)floor(v - 0.5)
j1 = j0 + 1
if (j0 < 0)
    j0 = h_t + j0
if (j1 > h_t - 1)
    j1 = j1 - h_t

```

```

w = (r - floor(r)) * d_t
k0 = (int)floor(w - 0.5)
k1 = k0 + 1
if (k0 < 0)
    k0 = d_t + k0
if (k1 > d_t - 1)
    k1 = k1 - d_t

```

```

a = frac(u - 0.5)
b = frac(v - 0.5)
c = frac(w - 0.5)

```

where $\text{frac}(x)$ denotes the fractional part of x and is computed as $x - \text{floor}(x)$.

For a 3D image, the image element value is found as

$$T = (1 - a) * (1 - b) * (1 - c) * T_{i_0j_0k_0}$$

$$\begin{aligned}
& + a * (1 - b) * (1 - c) * T_{i1j0k0} \\
& + (1 - a) * b * (1 - c) * T_{i0j1k0} \\
& + a * b * (1 - c) * T_{i1j1k0} \\
& + (1 - a) * (1 - b) * c * T_{i0j0k1} \\
& + a * (1 - b) * c * T_{i1j0k1} \\
& + (1 - a) * b * c * T_{i0j1k1} \\
& + a * b * c * T_{i1j1k1}
\end{aligned}$$

where T_{ijk} is the image element at location (i, j, k) in the 3D image.

For a 2D image, the image element value is found as

$$\begin{aligned}
T & = (1 - a) * (1 - b) * T_{i0j0} \\
& + a * (1 - b) * T_{i1j0} \\
& + (1 - a) * b * T_{i0j1} \\
& + a * b * T_{i1j1}
\end{aligned}$$

where T_{ij} is the image element at location (i, j) in the 2D image.

We now discuss how the addressing and filter modes are applied to generate the appropriate sample locations to read from the image if the addressing mode is `CLK_ADDRESS_MIRRORED_REPEAT`. The `CLK_ADDRESS_MIRRORED_REPEAT` addressing mode causes the image to be read as if it is tiled at every integer seam with the interpretation of the image data flipped at each integer crossing. For example, the (s, t, r) coordinates between 2 and 3 are addressed into the image as coordinates from 1 down to 0. If values in (s, t, r) are INF or NaN, the behavior of the built-in image read functions is undefined.

Filter Mode = CLK_FILTER_NEAREST

When filter mode is `CLK_FILTER_NEAREST`, the image element at location (i, j, k) becomes the image element value, with i, j and k computed as

$$\begin{aligned}
s' & = 2.0f * rint(0.5f * s) \\
s' & = fabs(s - s') \\
u & = s' * w_t \\
i & = (int) floor(u) \\
i & = min(i, w_t - 1) \\
\\
t' & = 2.0f * rint(0.5f * t) \\
t' & = fabs(t - t') \\
v & = t' * h_t \\
j & = (int) floor(v) \\
j & = min(j, h_t - 1)
\end{aligned}$$

```

r' = 2.0f * rint(0.5f * r)
r' = fabs(r - r')
w = r' * dt
k = (int)floor(w)
k = min(k, dt - 1)

```

For a 3D image, the image element at location (i, j, k) becomes the color value. For a 2D image, the image element at location (i, j) becomes the color value.

Filter Mode = CLK_FILTER_LINEAR

When filter mode is CLK_FILTER_LINEAR, a 2 × 2 square of image elements for a 2D image or a 2 × 2 × 2 cube of image elements for a 3D image is selected. This 2 × 2 square or 2 × 2 × 2 cube is obtained as follows.

Let

```

s' = 2.0f * rint(0.5f * s)
s' = fabs(s - s')
u = s' * wt
i0 = (int)floor(u - 0.5f)
i1 = i0 + 1
i0 = max(i0, 0)
i1 = min(i1, wt - 1)

t' = 2.0f * rint(0.5f * t)
t' = fabs(t - t')
v = t' * ht
j0 = (int)floor(v - 0.5f)
j1 = j0 + 1
j0 = max(j0, 0)
j1 = min(j1, ht - 1)

r' = 2.0f * rint(0.5f * r)
r' = fabs(r - r')
w = r' * dt
k0 = (int)floor(w - 0.5f)
k1 = k0 + 1
k0 = max(k0, 0)
k1 = min(k1, dt - 1)

a = frac(u - 0.5)
b = frac(v - 0.5)
c = frac(w - 0.5)

```

where $\text{frac}(x)$ denotes the fractional part of x and is computed as $x - \text{floor}(x)$.

For a 3D image, the image element value is found as

$$\begin{aligned} T = & (1 - a) * (1 - b) * (1 - c) * T_{i0j0k0} \\ & + a * (1 - b) * (1 - c) * T_{i1j0k0} \\ & + (1 - a) * b * (1 - c) * T_{i0j1k0} \\ & + a * b * (1 - c) * T_{i1j1k0} \\ & + (1 - a) * (1 - b) * c * T_{i0j0k1} \\ & + a * (1 - b) * c * T_{i1j0k1} \\ & + (1 - a) * b * c * T_{i0j1k1} \\ & + a * b * c * T_{i1j1k1} \end{aligned}$$

where T_{ijk} is the image element at location (i, j, k) in the 3D image.

For a 2D image, the image element value is found as

$$\begin{aligned} T = & (1 - a) * (1 - b) * T_{i0j0} \\ & + a * (1 - b) * T_{i1j0} \\ & + (1 - a) * b * T_{i0j1} \\ & + a * b * T_{i1j1} \end{aligned}$$

where T_{ij} is the image element at location (i, j) in the 2D image.

For a 1D image, the image element value is found as

$$\begin{aligned} T = & (1 - a) * T_{i0} \\ & + a * T_{i1} \end{aligned}$$

where T_i is the image element at location (i) in the 1D image.

NOTE

If the sampler is specified as using unnormalized coordinates (floating-point or integer coordinates), filter mode set to CLK_FILTER_NEAREST and addressing mode set to one of the following modes - CLK_ADDRESS_NONE, CLK_ADDRESS_CLAMP_TO_EDGE or CLK_ADDRESS_CLAMP, the location of the image element in the image given by (i, j, k) in section 8.2 will be computed without any loss of precision.

For all other sampler combinations of normalized or unnormalized coordinates, filter and addressing modes, the relative error or precision of the addressing mode calculations and the image filter operation are not defined by this revision of the OpenCL specification. To ensure a minimum precision of image addressing and filter calculations across any OpenCL device, for

these sampler combinations, developers should unnormalize the image coordinate in the kernel and implement the linear filter in the kernel with appropriate calls to `read_image{f|ui}` with a sampler that uses unnormalized coordinates, filter mode set to `CLK_FILTER_NEAREST`, addressing mode set to `CLK_ADDRESS_NONE`, `CLK_ADDRESS_CLAMP_TO_EDGE` or `CLK_ADDRESS_CLAMP` and finally performing the interpolation of color values read from the image to generate the filtered color value.

8.3 Conversion Rules

In this section we discuss conversion rules that are applied when reading and writing images in a kernel.

8.3.1 Conversion rules for normalized integer channel data types

In this section we discuss converting normalized integer channel data types to floating-point values and vice-versa.

8.3.1.1 Converting normalized integer channel data types to floating-point values

For images created with image channel data type of `CL_UNORM_INT8` and `CL_UNORM_INT16`, **read_imagef** will convert the channel values from an 8-bit or 16-bit unsigned integer to normalized floating-point values in the range $[0.0f \dots 1.0]$.

For images created with image channel data type of `CL_SNORM_INT8` and `CL_SNORM_INT16`, **read_imagef** will convert the channel values from an 8-bit or 16-bit signed integer to normalized floating-point values in the range $[-1.0 \dots 1.0]$.

These conversions are performed as follows:

`CL_UNORM_INT8` (8-bit unsigned integer) \rightarrow float

$$\text{normalized float value} = (\text{float})c / 255.0f$$

`CL_UNORM_INT_101010` (10-bit unsigned integer) \rightarrow float

$$\text{normalized float value} = (\text{float})c / 1023.0f$$

`CL_UNORM_INT16` (16-bit unsigned integer) \rightarrow float

$$\text{normalized float value} = (\text{float})c / 65535.0f$$

CL_SNORM_INT8 (8-bit signed integer) → float

normalized float value = $\max(-1.0f, (\text{float})c / 127.0f)$

CL_SNORM_INT16 (16-bit signed integer) → float

normalized float value = $\max(-1.0f, (\text{float})c / 32767.0f)$

The precision of the above conversions is ≤ 1.5 ulp except for the following cases.

For CL_UNORM_INT8

0 must convert to 0.0f and
255 must convert to 1.0f

For CL_UNORM_INT_101010

0 must convert to 0.0f and
1023 must convert to 1.0f

For CL_UNORM_INT16

0 must convert to 0.0f and
65535 must convert to 1.0f

For CL_SNORM_INT8

-128 and -127 must convert to -1.0f,
0 must convert to 0.0f and
127 must convert to 1.0f

For CL_SNORM_INT16

-32768 and -32767 must convert to -1.0f,
0 must convert to 0.0f and
32767 must convert to 1.0f

8.3.1.2 Converting floating-point values to normalized integer channel data types

For images created with image channel data type of CL_UNORM_INT8 and CL_UNORM_INT16, **write_imagef** will convert the floating-point color value to an 8-bit or 16-bit unsigned integer.

For images created with image channel data type of CL_SNORM_INT8 and CL_SNORM_INT16, **write_imagef** will convert the floating-point color value to an 8-bit or 16-bit signed integer.

The preferred method for how conversions from floating-point values to normalized integer values are performed is as follows:

float → CL_UNORM_INT8 (8-bit unsigned integer)

```
convert_uchar_sat_rte(f * 255.0f)
```

float → CL_UNORM_INT_101010 (10-bit unsigned integer)

```
min(convert_ushort_sat_rte(f * 1023.0f), 0x3ff)
```

float → CL_UNORM_INT16 (16-bit unsigned integer)

```
convert_ushort_sat_rte(f * 65535.0f)
```

float → CL_SNORM_INT8 (8-bit signed integer)

```
convert_char_sat_rte(f * 127.0f)
```

float → CL_SNORM_INT16 (16-bit signed integer)

```
convert_short_sat_rte(f * 32767.0f)
```

Please refer to *section 6.2.3.3* for out-of-range behavior and saturated conversions rules.

OpenCL implementations may choose to approximate the rounding mode used in the conversions described above. If a rounding mode other than round to nearest even (`_rte`) is used, the absolute error of the implementation dependant rounding mode vs. the result produced by the round to nearest even rounding mode must be ≤ 0.6 .

float → CL_UNORM_INT8 (8-bit unsigned integer)

```
Let  $f_{\text{preferred}} = \text{convert\_uchar\_sat\_rte}(f * 255.0f)$ 
```

```
Let  $f_{\text{approx}} =$   
    convert_uchar_sat_<impl-rounding-mode>(f * 255.0f)
```

```
 $\text{fabs}(f_{\text{preferred}} - f_{\text{approx}})$  must be  $\leq 0.6$ 
```

float → CL_UNORM_INT_101010 (10-bit unsigned integer)

```
Let  $f_{\text{preferred}} = \text{convert\_ushort\_sat\_rte}(f * 1023.0f)$ 
```

```
Let  $f_{\text{approx}} =$   
    convert_ushort_sat_<impl-rounding-mode>(f * 1023.0f)
```

```
 $\text{fabs}(f_{\text{preferred}} - f_{\text{approx}})$  must be  $\leq 0.6$ 
```

float → CL_UNORM_INT16 (16-bit unsigned integer)

```
Let fpreferred = convert_ushort_sat_rte(f * 65535.0f)
Let fapprox =
    convert_ushort_sat_<impl-rounding-mode>(f * 65535.0f)

fabs(fpreferred - fapprox) must be <= 0.6
```

float → CL_SNORM_INT8 (8-bit signed integer)

```
Let fpreferred = convert_char_sat_rte(f * 127.0f)
Let fapprox =
    convert_char_sat_<impl_rounding_mode>(f * 127.0f)

fabs(fpreferred - fapprox) must be <= 0.6
```

float → CL_SNORM_INT16 (16-bit signed integer)

```
Let fpreferred = convert_short_sat_rte(f * 32767.0f)
Let fapprox =
    convert_short_sat_<impl-rounding-mode>(f * 32767.0f)

fabs(fpreferred - fapprox) must be <= 0.6
```

8.3.2 Conversion rules for half precision floating-point channel data type

For images created with a channel data type of CL_HALF_FLOAT, the conversions from `half` to `float` are lossless (as described in *section 6.1.1.1*). Conversions from `float` to `half` round the mantissa using the round to nearest even or round to zero rounding mode. Denormalized numbers for the `half` data type which may be generated when converting a `float` to a `half` may be flushed to zero. A `float` NaN must be converted to an appropriate NaN in the `half` type. A `float` INF must be converted to an appropriate INF in the `half` type.

8.3.3 Conversion rules for floating-point channel data type

The following rules apply for reading and writing images created with channel data type of CL_FLOAT.

- ✚ NaNs may be converted to a NaN value(s) supported by the device.

- ✚ Denorms can be flushed to zero.
- ✚ All other values must be preserved.

8.3.4 Conversion rules for signed and unsigned 8-bit, 16-bit and 32-bit integer channel data types

Calls to **read_imagei** with channel data type values of CL_SIGNED_INT8, CL_SIGNED_INT16 and CL_SIGNED_INT32 return the unmodified integer values stored in the image at specified location.

Calls to **read_imageui** with channel data type values of CL_UNSIGNED_INT8, CL_UNSIGNED_INT16 and CL_UNSIGNED_INT32 return the unmodified integer values stored in the image at specified location.

Calls to **write_imagei** will perform one of the following conversions:

32 bit signed integer → 8-bit signed integer

```
convert_char_sat(i)
```

32 bit signed integer → 16-bit signed integer

```
convert_short_sat(i)
```

32 bit signed integer → 32-bit signed integer

```
no conversion is performed
```

Calls to **write_imageui** will perform one of the following conversions:

32 bit unsigned integer → 8-bit unsigned integer

```
convert_uchar_sat(i)
```

32 bit unsigned integer → 16-bit unsigned integer

```
convert_ushort_sat(i)
```

32 bit unsigned integer → 32-bit unsigned integer

```
no conversion is performed
```

The conversions described in this section must be correctly saturated.

8.4 Selecting an Image from an Image Array

Let (u, v, w) represent the unnormalized image coordinate values for reading from and/or writing to a 2D image in a 2D image array.

The 2D image layer selected is computed as:

$$\text{layer} = \text{clamp}(\text{rint}(w), 0, d_t - 1)$$

Let (u, v) represent the unnormalized image coordinate values for reading from and/or writing to a 1D image in a 1D image array.

The 1D image layer selected is computed as:

$$\text{layer} = \text{clamp}(\text{rint}(v), 0, h_t - 1)$$

9. Optional Extensions

The list of optional features supported by OpenCL 1.2 is described in the OpenCL 1.2 Extension Specification document.

10. OpenCL Embedded Profile

The OpenCL 1.2 specification describes the feature requirements for desktop platforms. This section describes the OpenCL 1.2 embedded profile that allows us to target a subset of the OpenCL 1.2 specification for handheld and embedded platforms. The optional extensions defined in the OpenCL 1.2 Extension Specification apply to both profiles.

The OpenCL 1.2 embedded profile has the following restrictions:

1. 64 bit integers i.e. long, along including the appropriate vector data types and operations on 64-bit integers are optional. The **cles_khr_int64**⁷³ extension string will be reported if the embedded profile implementation supports 64-bit integers.
2. Support for 3D images is optional.

If `CL_DEVICE_IMAGE3D_MAX_WIDTH`, `CL_DEVICE_IMAGE3D_MAX_HEIGHT` and `CL_DEVICE_IMAGE3D_MAX_DEPTH` are zero, the call to **clCreateImage** in the embedded profile will fail to create the 3D image. The *errcode_ret* argument in **clCreateImage** returns `CL_INVALID_OPERATION`. Declaring arguments of type `image3d_t` in a kernel will result in a compilation error.

If `CL_DEVICE_IMAGE3D_MAX_WIDTH`, `CL_DEVICE_IMAGE3D_HEIGHT` and `CL_DEVICE_IMAGE3D_MAX_DEPTH > 0`, 3D images are supported by the OpenCL embedded profile implementation. **clCreateImage** will work as defined by the OpenCL specification. The `image3d_t` data type can be used in a kernel(s).

3. Support for 2D image array writes is optional. If the **cles_khr_2d_image_array_writes** extension is supported by the embedded profile, writes to 2D image arrays are supported.
4. Image and image arrays created with an `image_channel_data_type` value of `CL_FLOAT` or `CL_HALF_FLOAT` can only be used with samplers that use a filter mode of `CL_FILTER_NEAREST`. The values returned by **read_imagef** and **read_imageh**⁷⁴ for 2D and 3D images if `image_channel_data_type` value is `CL_FLOAT` or `CL_HALF_FLOAT` and sampler with `filter_mode = CL_FILTER_LINEAR` are undefined.
5. The sampler addressing modes supported for image and image arrays are: `CLK_ADDRESS_NONE`, `CLK_ADDRESS_MIRRORED_REPEAT`, `CLK_ADDRESS_REPEAT`, `CLK_ADDRESS_CLAMP_TO_EDGE` and `CLK_ADDRESS_CLAMP`.

⁷³ Note that the performance of 64-bit integer arithmetic can vary significantly between embedded devices.

⁷⁴ If **cl_khr_fp16** extension is supported.

6. The mandated minimum single precision floating-point capability given by CL_DEVICE_SINGLE_FP_CONFIG is CL_FP_ROUND_TO_ZERO or CL_FP_ROUND_TO_NEAREST. If CL_FP_ROUND_TO_NEAREST is supported, the default rounding mode will be round to nearest even; otherwise the default rounding mode will be round to zero.
7. The single precision floating-point operations (addition, subtraction and multiplication) shall be correctly rounded. Zero results may always be positive 0.0. The accuracy of division and sqrt are given by *table 10.1*.

If CL_FP_INF_NAN is not set in CL_DEVICE_SINGLE_FP_CONFIG, and one of the operands or the result of addition, subtraction, multiplication or division would signal the overflow or invalid exception (see IEEE 754 specification), the value of the result is implementation-defined. Likewise, single precision comparison operators (<, >, <=, >=, ==, !=) return implementation-defined values when one or more operands is a NaN.

In all cases, conversions (*section 6.2* and *6.12.7*) shall be correctly rounded as described for the FULL_PROFILE, including those that consume or produce an INF or NaN. The built-in math functions (*section 6.12.2*) shall behave as described for the FULL_PROFILE, including edge case behavior described in *section 7.5.1* but with accuracy as described by *table 10.1*.

Note: If addition, subtraction and multiplication have default round to zero rounding mode, then **fract**, **fma** and **fdim** shall produce the correctly rounded result for round to zero rounding mode.

This relaxation of the requirement to adhere to IEEE 754 requirements for basic floating-point operations, though extremely undesirable, is to provide flexibility for embedded devices that have lot stricter requirements on hardware area budgets.

8. Denormalized numbers for the half data type which may be generated when converting a float to a half using variants of the **vstore_half** function or when converting from a half to a float using variants of the **vload_half** function can be flushed to zero. Refer to *section 6.1.1.1*.
9. The precision of conversions from CL_UNORM_INT8, CL_SNORM_INT8, CL_UNORM_INT16 and CL_SNORM_INT16 to float is ≤ 2 ulp for the embedded profile instead of ≤ 1.5 ulp as defined in *section 8.3.1.1*. The exception cases described in *section 8.3.1.1* and given below apply to the embedded profile.

For CL_UNORM_INT8

```
0 must convert to 0.0f and
255 must convert to 1.0f
```

For CL_UNORM_INT16

0 must convert to 0.0f and
65535 must convert to 1.0f

For CL_SNORM_INT8

-128 and -127 must convert to -1.0f,
0 must convert to 0.0f and
127 must convert to 1.0f

For CL_SNORM_INT16

-32768 and -32767 must convert to -1.0f,
0 must convert to 0.0f and
32767 must convert to 1.0f

For CL_UNORM_INT_101010

0 must convert to 0.0f and
1023 must convert to 1.0f

10. Built-in atomic functions as defined in *section 6.12.11* are optional.

The following optional extensions defined in the OpenCL 1.2 Extension Specification are available to the embedded profile:

- ✚ **cl_khr_int64_base_atomics**
- ✚ **cl_khr_int64_extended_atomics**
- ✚ **cl_khr_fp16**
- ✚ **cles_khr_int64**. If double precision is supported i.e. **CL_DEVICE_DOUBLE_FP_CONFIG** is not zero, then **cles_khr_int64** must also be supported.
- ✚ **cles_khr_2d_image_array_writes**. This extension indicates that the embedded profile device supports writes to 2D image arrays.

The following optional extensions defined in the OpenCL 1.0 and OpenCL 1.1 specifications (*section 9*) are also available to the embedded profile:

- ✚ **cl_khr_global_int32_base_atomics**
- ✚ **cl_khr_global_int32_extended_atomics**
- ✚ **cl_khr_local_int32_base_atomics**
- ✚ **cl_khr_local_int32_extended_atomics**

Table 10.1 describes the minimum accuracy of single precision floating-point arithmetic operations given as ULP values for the embedded profile. The reference value used to compute the ULP value of an arithmetic operation is the infinitely precise result.

| Function | Min Accuracy - ULP values ⁷⁵ |
|-----------------|---|
| $x + y$ | Correctly rounded |
| $x - y$ | Correctly rounded |
| $x * y$ | Correctly rounded |
| $1.0 / x$ | ≤ 3 ulp |
| x / y | ≤ 3 ulp |
| | |
| acos | ≤ 4 ulp |
| acospi | ≤ 5 ulp |
| asin | ≤ 4 ulp |
| asinpi | ≤ 5 ulp |
| atan | ≤ 5 ulp |
| atan2 | ≤ 6 ulp |
| atanpi | ≤ 5 ulp |
| atan2pi | ≤ 6 ulp |
| acosh | ≤ 4 ulp |
| asinh | ≤ 4 ulp |
| atanh | ≤ 5 ulp |
| cbrt | ≤ 4 ulp |
| ceil | Correctly rounded |
| copysign | 0 ulp |
| cos | ≤ 4 ulp |
| cosh | ≤ 4 ulp |
| cospi | ≤ 4 ulp |
| erfc | ≤ 16 ulp |
| erf | ≤ 16 ulp |
| exp | ≤ 4 ulp |
| exp2 | ≤ 4 ulp |
| exp10 | ≤ 4 ulp |
| expm1 | ≤ 4 ulp |
| fabs | 0 ulp |
| fdim | Correctly rounded |
| floor | Correctly rounded |
| fma | Correctly rounded |
| fmax | 0 ulp |
| fmin | 0 ulp |
| fmod | 0 ulp |
| fract | Correctly rounded |
| frexp | 0 ulp |
| hypot | ≤ 4 ulp |
| ilogb | 0 ulp |

⁷⁵ 0 ulp is used for math functions that do not require rounding.

| | |
|--------------------|-------------------------------------|
| ldexp | Correctly rounded |
| log | <= 4 ulp |
| log2 | <= 4 ulp |
| log10 | <= 4 ulp |
| log1p | <= 4 ulp |
| logb | 0 ulp |
| mad | Any value allowed (infinite ulp) |
| maxmag | 0 ulp |
| minmag | 0 ulp |
| modf | 0 ulp |
| nan | 0 ulp |
| nextafter | 0 ulp |
| pow(x, y) | <= 16 ulp |
| pown(x, y) | <= 16 ulp |
| powr(x, y) | <= 16 ulp |
| remainder | 0 ulp |
| remquo | 0 ulp |
| rint | Correctly rounded |
| rootn | <= 16 ulp |
| round | Correctly rounded |
| rsqrt | <= 4 ulp |
| sin | <= 4 ulp |
| sincos | <= 4 ulp for sine and cosine values |
| sinh | <= 4 ulp |
| sinpi | <= 4 ulp |
| sqrt | <= 4 ulp |
| tan | <= 5 ulp |
| tanh | <= 5 ulp |
| tanpi | <= 6 ulp |
| tgamma | <= 16 ulp |
| trunc | Correctly rounded |
| | |
| half_cos | <= 8192 ulp |
| half_divide | <= 8192 ulp |
| half_exp | <= 8192 ulp |
| half_exp2 | <= 8192 ulp |
| half_exp10 | <= 8192 ulp |
| half_log | <= 8192 ulp |
| half_log2 | <= 8192 ulp |
| half_log10 | <= 8192 ulp |
| half_powr | <= 8192 ulp |
| half_recip | <= 8192 ulp |
| half_rsqrt | <= 8192 ulp |
| half_sin | <= 8192 ulp |
| half_sqrt | <= 8192 ulp |

| | |
|----------------------|------------------------|
| half_tan | <= 8192 ulp |
| native_cos | Implementation-defined |
| native_divide | Implementation-defined |
| native_exp | Implementation-defined |
| native_exp2 | Implementation-defined |
| native_exp10 | Implementation-defined |
| native_log | Implementation-defined |
| native_log2 | Implementation-defined |
| native_log10 | Implementation-defined |
| native_powr | Implementation-defined |
| native_recip | Implementation-defined |
| native_rsqr | Implementation-defined |
| native_sin | Implementation-defined |
| native_sqrt | Implementation-defined |
| native_tan | Implementation-defined |

Table 10.1 *ULP values for built-in math functions*

The `__EMBEDDED_PROFILE__` macro is added to the language (refer to *section 6.10*). It will be the integer constant 1 for OpenCL devices that implement the embedded profile and is undefined otherwise.

`CL_PLATFORM_PROFILE` defined in *table 4.1* will return the string `EMBEDDED_PROFILE` if the OpenCL implementation supports the embedded profile only.

The minimum maximum values specified in *table 4.3* have been modified for the OpenCL embedded profile and are:

| cl_device_info | Return Type | Description |
|--|-----------------------------|--|
| <code>CL_DEVICE_TYPE</code> | <code>cl_device_type</code> | The OpenCL device type. Currently supported values are: <code>CL_DEVICE_TYPE_CPU</code> , <code>CL_DEVICE_TYPE_GPU</code> , <code>CL_DEVICE_TYPE_ACCELERATOR</code> , <code>CL_DEVICE_TYPE_DEFAULT</code> , a combination of the above types or <code>CL_DEVICE_TYPE_CUSTOM..</code> |
| <code>CL_DEVICE_VENDOR_ID</code> | <code>cl_uint</code> | A unique device vendor identifier. An example of a unique device identifier could be the PCIe ID. |
| <code>CL_DEVICE_MAX_COMPUTE_UNITS</code> | <code>cl_uint</code> | The number of parallel compute cores on the OpenCL device. The minimum value is 1. |

| | | |
|---|--------------|---|
| CL_DEVICE_MAX_WORK_ITEM_DIMENSIONS | unsigned int | Maximum dimensions that specify the global and local work-item IDs. The minimum value is 3 for devices that are not of type CL_DEVICE_TYPE_CUSTOM . |
| CL_DEVICE_MAX_WORK_ITEM_SIZES | size_t [] | Maximum number of work-items that can be specified in each dimension of the work-group to clEnqueueNDRangeKernel . Returns <i>n</i> size_t entries, where <i>n</i> is the value returned by the query for CL_DEVICE_MAX_WORK_ITEM_DIMENSIONS . The minimum value is (1, 1, 1) for devices that are not of type CL_DEVICE_TYPE_CUSTOM . |
| CL_DEVICE_MAX_WORK_GROUP_SIZE | size_t | Maximum number of work-items in a work-group executing a kernel using the data parallel execution model. (Refer to clEnqueueNDRangeKernel). The minimum value is 1. |
| CL_DEVICE_PREFERRED_VECTOR_WIDTH_CHAR CL_DEVICE_PREFERRED_VECTOR_WIDTH_SHORT CL_DEVICE_PREFERRED_VECTOR_WIDTH_INT CL_DEVICE_PREFERRED_VECTOR_WIDTH_LONG CL_DEVICE_PREFERRED_VECTOR_WIDTH_FLOAT CL_DEVICE_PREFERRED_VECTOR_WIDTH_DOUBLE CL_DEVICE_PREFERRED_VECTOR_WIDTH_HALF | cl_uint | Preferred native vector width size for built-in 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, CL_DEVICE_PREFERRED_VECTOR_WIDTH_DOUBLE must return 0. If the cl_khr_fp16 extension is not supported, CL_DEVICE_PREFERRED_VECTOR_WIDTH_HALF must return 0. |
| CL_DEVICE_NATIVE_VECTOR_WIDTH_CHAR CL_DEVICE_NATIVE_VECTOR_WIDTH_SHORT CL_DEVICE_NATIVE_VECTOR_WIDTH_INT CL_DEVICE_NATIVE_VECTOR_WIDTH_LONG | cl_uint | 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, CL_DEVICE_NATIVE_VECTOR_WIDTH_DOUBLE must return 0. |

| | | |
|---|--------------------|--|
| CL_DEVICE_NATIVE_VECTOR_WIDTH_FLOAT CL_DEVICE_NATIVE_VECTOR_WIDTH_DOUBLE CL_DEVICE_NATIVE_VECTOR_WIDTH_HALF | | If the cl_khr_fp16 extension is not supported, CL_DEVICE_NATIVE_VECTOR_WIDTH_HALF must return 0. |
| CL_DEVICE_MAX_CLOCK_FREQUENCY | cl_uint | Maximum configured clock frequency of the device in MHz. |
| CL_DEVICE_ADDRESS_BITS | cl_uint | The default compute device address space size specified as an unsigned integer value in bits. Currently supported values are 32 or 64 bits. If the value reported by the embedded profile is 64, then the cles_khr_int64 extension must be supported. |
| CL_DEVICE_MAX_WORK_GROUP_SIZE | size_t | Maximum number of work-items in a work-group executing a kernel using the data parallel execution model. (Refer to clEnqueueNDRangeKernel). The minimum value is 1. |
| | | |
| CL_DEVICE_MAX_MEM_ALLOC_SIZE | unsigned long long | Max size of memory object allocation in bytes. The minimum value is max (1/4 th of CL_DEVICE_GLOBAL_MEM_SIZE , 1*1024*1024) for devices that are not of type CL_DEVICE_TYPE_CUSTOM. |
| | | |
| CL_DEVICE_IMAGE_SUPPORT | cl_bool | Is CL_TRUE if images are supported by the OpenCL device and CL_FALSE otherwise. |
| CL_DEVICE_MAX_READ_IMAGE_ARGS | unsigned int | Max number of simultaneous image objects that can be read by a kernel. The minimum value is 8 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_MAX_WRITE_IMAGE_ARGS | unsigned int | Max number of simultaneous image objects that can be written to by a kernel. The minimum value is 1 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_IMAGE2D_MAX_WIDTH | size_t | Max width of 2D image in pixels. The minimum value is 2048 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |

| | | |
|--|--------------|--|
| CL_DEVICE_IMAGE2D_MAX_HEIGHT | size_t | Max height of 2D image in pixels. The minimum value is 2048 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_IMAGE3D_MAX_WIDTH | size_t | Max width of 3D image in pixels. The minimum value is 0 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_IMAGE3D_MAX_HEIGHT | size_t | Max height of 3D image in pixels. The minimum value is 0. |
| CL_DEVICE_IMAGE3D_MAX_DEPTH | size_t | Max depth of 3D image in pixels. The minimum value is 0. |
| CL_DEVICE_IMAGE_MAX_BUFFER_SIZE | size_t | Max number of pixels for a 1D image created from a buffer object. The minimum value is 2048 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_IMAGE_MAX_ARRAY_SIZE | size_t | Max number of images in a 1D or 2D image array. The minimum value is 256 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| CL_DEVICE_MAX_SAMPLERS | unsigned int | Maximum number of samplers that can be used in a kernel. Refer to <i>section 6.12.14</i> for a detailed description on samplers. The minimum value is 8 if CL_DEVICE_IMAGE_SUPPORT is CL_TRUE. |
| | | |
| CL_DEVICE_MAX_PARAMETER_SIZE | size_t | Max size in bytes of the arguments that can be passed to a kernel. The minimum value is 256 bytes for devices that are not of type CL_DEVICE_TYPE_CUSTOM. |
| | | |
| CL_DEVICE_MEM_BASE_ADDR_ALIGN | cl_uint | The minimum value is the size (in bits) of the largest OpenCL built-in data type supported by the device (long16 in FULL profile, long16 or int16 in EMBEDDED profile) for devices that are not of type CL_DEVICE_TYPE_CUSTOM. |

| | | |
|--|----------------------------|--|
| <p>CL_DEVICE_SINGLE_FP_CONFIG</p> | <p>cl_device_fp_config</p> | <p>Describes single precision floating-point capability of the device. This is a bit-field that describes one or more of the following values:</p> <p>CL_FP_DENORM – denorms are supported</p> <p>CL_FP_INF_NAN – INF and quiet NaNs are supported.</p> <p>CL_FP_ROUND_TO_NEAREST– round to nearest even rounding mode supported</p> <p>CL_FP_ROUND_TO_ZERO – round to zero rounding mode supported</p> <p>CL_FP_ROUND_TO_INF – round to positive and negative infinity rounding modes supported</p> <p>CL_FP_FMA – IEEE754-2008 fused multiply-add is supported.</p> <p>CL_FP_CORRECTLY_ROUNDED_DIVIDE_SQRT – divide and sqrt are correctly rounded as defined by the IEEE754 specification.</p> <p>CL_FP_SOFT_FLOAT – Basic floating-point operations (such as addition, subtraction, multiplication) are implemented in software.</p> <p>The mandated minimum floating-point capability is: CL_FP_ROUND_TO_ZERO or CL_FP_ROUND_TO_NEAREST for devices that are not of type CL_DEVICE_TYPE_CUSTOM.</p> |
| <p>CL_DEVICE_DOUBLE_FP_CONFIG</p> | <p>cl_device_fp_config</p> | <p>Describes double precision floating-point capability of the OpenCL device. This is a bit-field that describes one or more of the following values:</p> <p>CL_FP_DENORM – denorms are supported</p> <p>CL_FP_INF_NAN – INF and NaNs are supported.</p> <p>CL_FP_ROUND_TO_NEAREST – round to nearest even rounding mode supported.</p> <p>CL_FP_ROUND_TO_ZERO – round to zero</p> |

| | | |
|--|--------------------------|---|
| | | <p>rounding mode supported.</p> <p>CL_FP_ROUND_TO_INF – round to positive and negative infinity rounding modes supported.</p> <p>CP_FP_FMA – IEEE754-2008 fused multiply-add is supported.</p> <p>CL_FP_SOFT_FLOAT – Basic floating-point operations (such as addition, subtraction, multiplication) are implemented in software.</p> <p>Double precision is an optional feature so the mandated minimum double precision floating-point capability is 0.</p> <p>If double precision is supported by the device, then the minimum double precision floating-point capability must be:</p> <p>CL_FP_FMA CL_FP_ROUND_TO_NEAREST CL_FP_ROUND_TO_ZERO CL_FP_ROUND_TO_INF CL_FP_INF_NAN CL_FP_DENORM.</p> |
| CL_DEVICE_GLOBAL_MEM_CACHE_TYPE | cl_device_mem_cache_type | Type of global memory cache supported. Valid values are: CL_NONE, CL_READ_ONLY_CACHE and CL_READ_WRITE_CACHE. |
| CL_DEVICE_GLOBAL_MEM_CACHELINE_SIZE | cl_uint | Size of global memory cache line in bytes. |
| CL_DEVICE_GLOBAL_MEM_CACHE_SIZE | cl_ulong | Size of global memory cache in bytes. |
| CL_DEVICE_GLOBAL_MEM_SIZE | cl_ulong | Size of global device memory in bytes. |
| | | |
| CL_DEVICE_MAX_CONSTANT_BUFFER_SIZE | unsigned long long | Max size in bytes of a constant buffer allocation. The minimum value is 1 KB for devices that are not of type CL_DEVICE_TYPE_CUSTOM. |
| CL_DEVICE_MAX_CONSTANT_ARGS | unsigned int | Max number of arguments declared with the __constant qualifier in a kernel. The minimum value is 4 for devices that are not of type CL_DEVICE_TYPE_CUSTOM. |
| | | |
| CL_DEVICE_LOCAL_MEM_TYPE | cl_device_local_mem_type | Type of local memory supported. This can be set to CL_LOCAL implying |

| | | |
|---|----------|--|
| | | dedicated local memory storage such as SRAM, or CL_GLOBAL. For custom devices, CL_NONE can also be returned indicating no local memory support. |
| CL_DEVICE_LOCAL_MEM_SIZE | cl_ulong | Size of local memory arena in bytes. The minimum value is 1 KB for devices that are not of type CL_DEVICE_TYPE_CUSTOM. |
| CL_DEVICE_ERROR_CORRECTION_SUPPORT | cl_bool | Is CL_TRUE if the device implements error correction for all accesses to compute device memory (global and constant). Is CL_FALSE if the device does not implement such error correction. |
| | | |
| CL_DEVICE_HOST_UNIFIED_MEMORY | cl_bool | Is CL_TRUE if the device and the host have a unified memory subsystem and is CL_FALSE otherwise. |
| | | |
| CL_DEVICE_PROFILING_TIMER_RESOLUTION | size_t | Describes the resolution of device timer. This is measured in nanoseconds. Refer to <i>section 5.12</i> for details. |
| | | |
| CL_DEVICE_ENDIAN_LITTLE | cl_bool | Is CL_TRUE if the OpenCL device is a little endian device and CL_FALSE otherwise. |
| CL_DEVICE_AVAILABLE | cl_bool | Is CL_TRUE if the device is available and CL_FALSE if the device is not available. |
| | | |
| CL_DEVICE_COMPILER_AVAILABLE | cl_bool | Is CL_FALSE if the implementation does not have a compiler available to compile the program source. Is CL_TRUE if the compiler is available. This can be CL_FALSE for the embedded platform profile only. |
| CL_DEVICE_LINKER_AVAILABLE | cl_bool | Is CL_FALSE if the implementation does not have a linker available. Is CL_TRUE if the linker is available. |

| | | |
|---|-----------------------------|--|
| | | <p>This can be CL_FALSE for the embedded platform profile only.</p> <p>This must be CL_TRUE if CL_DEVICE_COMPILER_AVAILABLE is CL_TRUE.</p> |
| CL_DEVICE_EXECUTION_CAPABILITIES | cl_device_exec_capabilities | <p>Describes the execution capabilities of the device. This is a bit-field that describes one or more of the following values:</p> <p>CL_EXEC_KERNEL – The OpenCL device can execute OpenCL kernels.</p> <p>CL_EXEC_NATIVE_KERNEL – The OpenCL device can execute native kernels.</p> <p>The mandated minimum capability is: CL_EXEC_KERNEL.</p> |
| CL_DEVICE_QUEUE_PROPERTIES | cl_command_queue_properties | <p>Describes the command-queue properties supported of the device. This is a bit-field that describes one or more of the following values:</p> <p>CL_QUEUE_OUT_OF_ORDER_EXEC_MODE_ENABLE</p> <p>CL_QUEUE_PROFILING_ENABLE</p> <p>These properties are described in <i>table 5.1</i>.</p> <p>The mandated minimum capability is: CL_QUEUE_PROFILING_ENABLE.</p> |
| CL_DEVICE_BUILT_IN_KERNELS | char[] | <p>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.</p> |
| CL_DEVICE_PRINTF_BUFFER_SIZE | size_t | <p>Maximum size of the internal buffer that holds the output of printf calls from a kernel. The minimum value for the EMBEDDED profile is 1 KB.</p> |

| | | |
|--|--------------------------------|--|
| CL_DEVICE_PREFERRED_INTEROP_USER_SYNC | cl_bool | Is CL_TRUE if the device's preference is for the user to be responsible for synchronization, when sharing memory objects between OpenCL and other APIs such as DirectX, CL_FALSE if the device / implementation has a performant path for performing synchronization of memory object shared between OpenCL and other APIs such as DirectX. |
| CL_DEVICE_PARENT_DEVICE | cl_device_id | Returns the cl_device_id of the parent device to which this sub-device belongs. If <i>device</i> is a root-level device, a NULL value is returned. |
| CL_DEVICE_PARTITION_MAX_SUB_DEVICES | cl_uint | Returns the maximum number of sub-devices that can be created when a device is partitioned. The value returned cannot exceed CL_DEVICE_MAX_COMPUTE_UNITS. |
| CL_DEVICE_PARTITION_PROPERTIES | cl_device_partition_property[] | Returns the list of partition types supported by <i>device</i> . The return type is an array of cl_device_partition_property values drawn from the following list: CL_DEVICE_PARTITION_EQUALLY CL_DEVICE_PARTITION_BY_COUNTS CL_DEVICE_PARTITION_BY_AFFINITY_DOMAIN If the device does not support any partition types, a value of 0 will be returned. |
| CL_DEVICE_PARTITION_AFFINITY_DOMAIN | cl_device_affinity_domain | Returns the list of supported affinity domains for partitioning the <i>device</i> using CL_DEVICE_PARTITION_BY_AFFINITY_DOMAIN. This is a bit-field that describes one or more of the following values: CL_AFFINITY_DOMAIN_NUMA CL_AFFINITY_DOMAIN_L4_CACHE CL_AFFINITY_DOMAIN_L3_CACHE CL_AFFINITY_DOMAIN_L2_CACHE CL_AFFINITY_DOMAIN_L1_CACHE |

| | | |
|---------------------------|--------------------------------|---|
| | | CL_AFFINITY_DOMAIN_NEXT_PARTITIONABLE If the device does not support any affinity domains, a value of 0 will be returned. |
| CL_DEVICE_PARTITION_TYPE | cl_device_partition_property[] | Returns the properties argument specified in clCreateSubDevices if <i>device</i> is a sub-device. Otherwise the implementation may either return a <i>param_value_size_ret</i> 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 <i>param_value</i> points to. |
| CL_DEVICE_REFERENCE_COUNT | cl_uint | Returns the <i>device</i> reference count. If the device is a root-level device, a reference count of one is returned. |

If CL_DEVICE_IMAGE_SUPPORT specified in *table 4.3* is CL_TRUE, the values assigned to CL_DEVICE_MAX_READ_IMAGE_ARGS, CL_DEVICE_MAX_WRITE_IMAGE_ARGS, CL_DEVICE_IMAGE2D_MAX_WIDTH, CL_DEVICE_IMAGE2D_MAX_HEIGHT, CL_DEVICE_IMAGE3D_MAX_WIDTH, CL_DEVICE_IMAGE3D_MAX_HEIGHT, CL_DEVICE_IMAGE3D_MAX_DEPTH and CL_DEVICE_MAX_SAMPLERS by the implementation must be greater than or equal to the minimum values specified in the embedded profile version of *table 4.3* given above. In addition, the following list of image formats must be supported by the OpenCL embedded profile implementation.

For 1D, 2D, optional 3D images, 1D and 2D image array objects, the minimum list of supported image formats (for reading and writing) is:

| image_num_channels | image_channel_order | image_channel_data_type |
|--------------------|---------------------|--|
| 4 | CL_RGBA | CL_UNORM_INT8 CL_UNORM_INT16 CL_SIGNED_INT8 CL_SIGNED_INT16 CL_SIGNED_INT32 CL_UNSIGNED_INT8 CL_UNSIGNED_INT16 CL_UNSIGNED_INT32 CL_HALF_FLOAT CL_FLOAT |

11. References

1. The ISO/IEC 9899:1999 “C” Language Specification.
2. The ISO/IEC JTC1 SC22 WG14 N1169 Specification.
3. The ANSI/IEEE Std 754-1985 and 754-2008 Specifications.
4. The AltiVec™ Technology Programming Interface Manual.
5. Ian Buck, Tim Foley, Daniel Horn, Jeremy Sugarman, Kayvon Fatahalian, Mike Houston, Pat Hanrahan. *Brook for GPUs: Stream Computing on Graphics Hardware*
6. Ian Buck. *Brook Specification v0.2*.
<http://merrimac.stanford.edu/brook/brookspec-v0.2.pdf>
7. *NVIDIA CUDA Programming Guide*.
<http://developer.nvidia.com/object/cuda.html>
8. *ATI CTM Guide – Technical Reference Manual*
http://ati.amd.com/companyinfo/researcher/documents/ATI_CTM_Guide.pdf
9. *OpenMP Application Program Interface*.
<http://www.openmp.org/drupal/mp-documents/spec25.pdf>
10. *The OpenGL Specification and the OpenGL Shading Language Specification*
<http://www.opengl.org/registry/>
11. *NESL – A nested data parallel language*.
<http://www.cs.cmu.edu/~scandal/nesl.html>
12. ***On the definition of ulp*** (x) by Jean-Michel Muller
<ftp://ftp.inria.fr/INRIA/publication/publi-pdf/RR/RR-5504.pdf>
13. Explicit Memory Fences
<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2007/n2262.html>
14. Lefohn, Kniss, Strzodka, Sengupta, Owens, "Glift: Generic, Efficient, Random-Access GPU Data Structures," ACM Transactions on Graphics, Jan. 2006. pp 60--99.
15. Pharr, Lefohn, Kolb, Lalonde, Foley, Berry, "Programmable Graphics---The Future of Interactive Rendering," Neoptica Whitepaper, Mar. 2007.
16. Jens Maurer, Michael Wong. Towards support for attributes in C++ (Revision 4). March 2008. Proposed to WG21 “Programming Language C++, Core Working Group”.
17. GCC Attribute Syntax. <http://gcc.gnu.org/onlinedocs/gcc/Attribute-Syntax.html>.

Appendix A

A.1 Shared OpenCL Objects

This section describes which objects can be shared across multiple command-queues created within a host process.

OpenCL memory objects, program objects and kernel objects are created using a context and can be shared across multiple command-queues created using the same context. Event objects can be created when a command is queued to a command-queue. These event objects can be shared across multiple command-queues created using the same context.

The application needs to implement appropriate synchronization across threads on the host processor to ensure that the changes to the state of a shared object (such as a command-queue object, memory object, program or kernel object) happen in the correct order (deemed correct by the application) when multiple command-queues in multiple threads are making changes to the state of a shared object.

A command-queue can cache changes to the state of a memory object on the device associated with the command-queue. To synchronize changes to a memory object across command-queues, the application must do the following:

In the command-queue that includes commands that modify the state of a memory object, the application must do the following:

- ✚ Get appropriate event objects for commands that modify the state of the shared memory object.
- ✚ Call the **clFlush** (or **clFinish**) API to issue any outstanding commands from this command-queue.

In the command-queue that wants to synchronize to the latest state of a memory object, commands queued by the application must use the appropriate event objects that represent commands that modify the state of the shared memory object as event objects to wait on. This is to ensure that commands that use this shared memory object complete in the previous command-queue before the memory objects are used by commands executing in this command-queue.

The results of modifying a shared resource in one command-queue while it is being used by another command-queue are undefined.

A.2 Multiple Host Threads

All OpenCL API calls are thread-safe⁷⁶ except **clSetKernelArg**. **clSetKernelArg** is safe to call from any host thread, and is safe to call re-entrantly so long as concurrent calls operate on different `cl_kernel` objects. However, the behavior of the `cl_kernel` object is undefined if **clSetKernelArg** is called from multiple host threads on the same `cl_kernel` object at the same time⁷⁷. Please note that there are additional limitations as to which OpenCL APIs may be called from OpenCL callback functions -- please see *section 5.9*.

The behavior of OpenCL APIs called from an interrupt or signal handler is implementation-defined

The OpenCL implementation should be able to create multiple command-queues for a given OpenCL context and multiple OpenCL contexts in an application running on the host processor.

⁷⁶ Please refer to the OpenCL glossary for the OpenCL definition of thread-safe. This definition may be different from usage of the term in other contexts.

⁷⁷ There is an inherent race condition in the design of OpenCL that occurs between setting a kernel argument and using the kernel with `clEnqueueNDRangeKernel` or `clEnqueueTask`. Another host thread might change the kernel arguments between when a host thread sets the kernel arguments and then enqueues the kernel, causing the wrong kernel arguments to be enqueued. Rather than attempt to share `cl_kernel` objects among multiple host threads, applications are strongly encouraged to make additional `cl_kernel` objects for kernel functions for each host thread.

Appendix B — Portability

OpenCL is designed to be portable to other architectures and hardware designs. OpenCL uses at its core a C99 based programming language. Floating-point arithmetic is based on the **IEEE-754** and **IEEE-754-2008** standards. The memory objects, pointer qualifiers and weakly ordered memory are designed to provide maximum compatibility with discrete memory architectures implemented by OpenCL devices. Command-queues and barriers allow for synchronization between the host and OpenCL devices. The design, capabilities and limitations of OpenCL are very much a reflection of the capabilities of underlying hardware.

Unfortunately, there are a number of areas where idiosyncrasies of one hardware platform may allow it to do some things that do not work on another. By virtue of the rich operating system resident on the CPU, on some implementations the kernels executing on a CPU may be able to call out to system services whereas the same calls on the GPU will likely fail for now. (Please see *section 6.9*). Since there is some advantage to having these services available for debugging purposes, implementations can use the OpenCL extension mechanism to implement these services.

Likewise, the heterogeneity of computing architectures might mean that a particular loop construct might execute at an acceptable speed on the CPU but very poorly on a GPU, for example. CPUs are designed in general to work well on latency sensitive algorithms on single threaded tasks, whereas common GPUs may encounter extremely long latencies, potentially orders of magnitude worse. A developer interested in writing portable code may find that it is necessary to test his design on a diversity of hardware designs to make sure that key algorithms are structured in a way that works well on a diversity of hardware. We suggest favoring more work-items over fewer. It is anticipated that over the coming months and years experience will produce a set of best practices that will help foster a uniformly favorable experience on a diversity of computing devices.

Of somewhat more concern is the topic of endianness. Since a majority of devices supported by the initial implementation of OpenCL are little-endian, developers need to make sure that their kernels are tested on both big-endian and little-endian devices to ensure source compatibility with OpenCL devices now and in the future. The endian attribute qualifier is supported by the OpenCL C programming language to allow developers to specify whether the data uses the endianness of the host or the OpenCL device. This allows the OpenCL compiler to do appropriate endian-conversion on load and store operations from or to this data.

We also describe how endianness can leak into an implementation causing kernels to produce unintended results:

When a big-endian vector machine (e.g. AltiVec, CELL SPE) loads a vector, the order of the data is retained. That is both the order of the bytes within each element and the order of the elements in the vector are the same as in memory. When a little-endian vector machine (e.g. SSE) loads a vector, the order of the data in register (where all the work is done) is reversed. **Both** the order of the bytes within each element and the order of the elements with respect to one

another in the vector are reversed.

Memory:

| | | | | |
|-----------|-------------------|------------|------------|------------|
| uint4 a = | 0x00010203 | 0x04050607 | 0x08090A0B | 0x0C0D0E0F |
|-----------|-------------------|------------|------------|------------|

In register (big-endian):

| | | | | |
|-----------|-------------------|------------|------------|------------|
| uint4 a = | 0x00010203 | 0x04050607 | 0x08090A0B | 0x0C0D0E0F |
|-----------|-------------------|------------|------------|------------|

In register (little-endian):

| | | | | |
|-----------|------------|------------|------------|-------------------|
| uint4 a = | 0x0F0E0D0C | 0x0B0A0908 | 0x07060504 | 0x03020100 |
|-----------|------------|------------|------------|-------------------|

This allows little-endian machines to use a single vector load to load little-endian data, regardless of how large each piece of data is in the vector. That is the transformation is equally valid whether that vector was a `uchar16` or a `ulong2`. Of course, as is well known, little-endian machines actually⁷⁸ store their data in reverse byte order to compensate for the little-endian storage format of the array elements:

Memory (big-endian):

| | | | | |
|-----------|-------------------|------------|------------|------------|
| uint4 a = | 0x00010203 | 0x04050607 | 0x08090A0B | 0x0C0D0E0F |
|-----------|-------------------|------------|------------|------------|

Memory (little-endian):

| | | | | |
|-----------|-------------------|------------|------------|------------|
| uint4 a = | 0x03020100 | 0x07060504 | 0x0B0A0908 | 0x0F0E0D0C |
|-----------|-------------------|------------|------------|------------|

Once that data is loaded into a vector, we end up with this:

In register (big-endian):

| | | | | |
|-----------|-------------------|------------|------------|------------|
| uint4 a = | 0x00010203 | 0x04050607 | 0x08090A0B | 0x0C0D0E0F |
|-----------|-------------------|------------|------------|------------|

In register (little-endian):

| | | | | |
|-----------|------------|------------|------------|-------------------|
| uint4 a = | 0x0C0D0E0F | 0x08090A0B | 0x04050607 | 0x00010203 |
|-----------|------------|------------|------------|-------------------|

⁷⁸ Note that we are talking about the programming model here. In reality, little endian systems might choose to simply address their bytes from "the right" or reverse the "order" of the bits in the byte. Either of these choices would mean that no big swap would need to occur in hardware.

That is, in the process of correcting the endianness of the bytes within each element, the machine ends up reversing the order that the elements appear in the vector with respect to each other within the vector. `0x00010203` appears at the left of the big-endian vector and at the right of the little-endian vector.

When the host and device have different endianness, the developer must ensure that kernel argument values are processed correctly. The implementation may or may not automatically convert endianness of kernel arguments. Developers should consult vendor documentation for guidance on how to handle kernel arguments in these situations.

OpenCL provides a consistent programming model across architectures by numbering elements according to their order in memory. Concepts such as *even/odd* and *high/low* follow accordingly. Once the data is loaded into registers, we find that element 0 is at the left of the big-endian vector and element 0 is at the right of the little-endian vector:

```
float x[4];
float4 v = vload4( 0, x );
```

Big-endian:

```
v contains { x[0], x[1], x[2], x[3] }
```

Little-endian:

```
v contains { x[3], x[2], x[1], x[0] }
```

The compiler is aware that this swap occurs and references elements accordingly. So long as we refer to them by a numeric index such as `.s0123456789abcdef` or by descriptors such as `.xyzw`, `.hi`, `.lo`, `.even` and `.odd`, everything works transparently. Any ordering reversal is undone when the data is stored back to memory. The developer should be able to work with a big endian programming model and ignore the element ordering problem in the vector ... for most problems. This mechanism relies on the fact that we can rely on a consistent element numbering. Once we change numbering system, for example by conversion-free casting (using `as_typen`) a vector to another vector of the same size but a different number of elements, then we get different results on different implementations depending on whether the system is big- endian, or little-endian or indeed has no vector unit at all. (Thus, the behavior of bitcasts to vectors of different numbers of elements is implementation-defined, see *section 6.2.4*)

An example follows:

```
float x[4] = { 0.0f, 1.0f, 2.0f, 3.0f };
float4 v = vload4( 0, x );
uint4 y = (uint4) v;           // legal, portable
ushort8 z = (ushort8) v;      // legal, not portable
                                // element size changed
```

Big-endian:

```

v contains { 0.0f, 1.0f, 2.0f, 3.0f }
y contains { 0x00000000, 0x3f800000,
             0x40000000, 0x40400000 }
z contains { 0x0000, 0x0000, 0x3f80, 0x0000,
             0x4000, 0x0000, 0x4040, 0x0000 }
z.z is 0x3f80

```

Little-endian:

```

v contains { 3.0f, 2.0f, 1.0f, 0.0f }
y contains { 0x40400000, 0x40000000,
             0x3f800000, 0x00000000 }
z contains { 0x4040, 0x0000, 0x4000,
             0x0000, 0x3f80, 0x0000, 0x0000, 0x0000
             }
z.z is 0

```

Here, the value in `z.z` is not the same between big- and little-endian vector machines

OpenCL could have made it illegal to do a conversion free cast that changes the number of elements in the name of portability. However, while OpenCL provides a common set of operators drawing from the set that are typically found on vector machines, it can not provide access to everything every ISA may offer in a consistent uniform portable manner. Many vector ISAs provide special purpose instructions that greatly accelerate specific operations such as DCT, SAD, or 3D geometry. It is not intended for OpenCL to be so heavy handed that time-critical performance sensitive algorithms can not be written by knowledgeable developers to perform at near peak performance. Developers willing to throw away portability should be able to use the platform-specific instructions in their code. For this reason, OpenCL is designed to allow traditional vector C language programming extensions, such as the AltiVec C Programming Interface or the Intel C programming interfaces (such as those found in `emmintrin.h`) to be used directly in OpenCL with OpenCL data types as an extension to OpenCL. As these interfaces rely on the ability to do conversion-free casts that change the number of elements in the vector to function properly, OpenCL allows them too.

As a general rule, any operation that operates on vector types in segments that are not the same size as the vector element size may break on other hardware with different endianness or different vector architecture.

Examples might include:

- ✚ Combining two `uchar8`'s containing high and low bytes of a `ushort`, to make a `ushort8` using `.even` and `.odd` operators (please use **`upsample()`** for this, see *section 6.12.3*)
- ✚ Any bitcast that changes the number of elements in the vector. (Operations on the new type are non-portable.)

- ✚ Swizzle operations that change the order of data using chunk sizes that are not the same as the element size

Examples of operations that are portable:

- ✚ Combining two `uint8`'s to make a `uchar16` using `.even` and `.odd` operators. For example to interleave left and right audio streams.
- ✚ Any bitcast that does not change the number of elements (e.g. `(float4) uint4` -- we define the storage format for floating-point types)
- ✚ Swizzle operations that swizzle elements of the same size as the elements of the vector.

OpenCL has made some additions to C to make application behavior more dependable than C. Most notably in a few cases OpenCL defines the behavior of some operations that are undefined in C99:

- ✚ OpenCL provides `convert_` operators for conversion between all types. C99 does not define what happens when a floating-point type is converted to integer type and the floating-point value lies outside the representable range of the integer type after rounding. When the `_sat` variant of the conversion is used, the float shall be converted to the nearest representable integer value. Similarly, OpenCL also makes recommendations about what should happen with NaN. Hardware manufacturers that provide the saturated conversion in hardware may use the saturated conversion hardware for both the saturated and non-saturated versions of the OpenCL `convert_` operator. OpenCL does not define what happens for the non-saturated conversions when floating-point operands are outside the range representable integers after rounding.
- ✚ The format of `half`, `float`, and `double` types is defined to be the `binary16`, `binary32` and `binary64` formats in the draft IEEE-754 standard. (The latter two are identical to the existing IEEE-754 standard.) You may depend on the positioning and meaning of the bits in these types.
- ✚ OpenCL defines behavior for oversized shift values. Shift operations that shift greater than or equal to the number of bits in the first operand reduce the shift value modulo the number of bits in the element. For example, if we shift an `int4` left by 33 bits, OpenCL treats this as shift left by $33\%32 = 1$ bit.
- ✚ A number of edge cases for math library functions are more rigorously defined than in C99. Please see *section 7.5*.

Appendix C — Application Data Types

This section documents the provided host application types and constant definitions. The documented material describes the commonly defined data structures, types and constant values available to all platforms and architectures. The addition of these details demonstrates our commitment to maintaining a portable programming environment and potentially deters changes to the supplied headers.

C.1 Shared Application Scalar Data Types

The following application scalar types are provided for application convenience.

```
cl_char
cl_uchar
cl_short
cl_ushort
cl_int
cl_uint
cl_long
cl_ulong
cl_half
cl_float
cl_double
```

C.2 Supported Application Vector Data Types

Application vector types are unions used to create vectors of the above application scalar types. The following application vector types are provided for application convenience.

```
cl_charn
cl_ucharn
cl_shortn
cl_ushortn
cl_intn
cl_uintn
cl_longn
cl_ulongn
cl_halfn
cl_floatn
cl_doublen
```

n can be 2, 3, 4, 8 or 16.

The application scalar and vector data types are defined in the `cl_platform.h` header file.

C.3 Alignment of Application Data Types

The user is responsible for ensuring that data passed into and out of OpenCL buffers are natively aligned relative to the start of the buffer per requirements in *section 6.1.5*. This implies that OpenCL buffers created with `CL_MEM_USE_HOST_PTR` need to provide an appropriately aligned host memory pointer that is aligned to the data types used to access these buffers in a kernel(s). As well, the user is responsible to ensure that data passed into and out of OpenCL images are properly aligned to the granularity of the data representing a single pixel (e.g. `image_num_channels * sizeof(image_channel_data_type)`) except for `CL_RGB` and `CL_RGBx` images where the data must be aligned to the granularity of a single channel in a pixel (i.e. `sizeof(image_channel_data_type)`).

OpenCL makes no requirement about the alignment of OpenCL application defined data types outside of buffers and images, except that the underlying vector primitives (e.g. `__cl_float4`) where defined shall be directly accessible as such using appropriate named fields in the `cl_type` union (see *section C.5*). Nevertheless, it is recommended that the `cl_platform.h` header should attempt to naturally align OpenCL defined application data types (e.g. `cl_float4`) according to their type.

C.4 Vector Literals

Application vector literals may be used in assignments of individual vector components. Literal usage follows the convention of the underlying application compiler.

```
cl_float2 foo = { .s[1] = 2.0f };
cl_int8 bar = {{ 2, 4, 6, 8, 10, 12, 14, 16 }};
```

C.5 Vector Components

The components of application vector types can be addressed using the `<vector_name>.s[<index>]` notation.

For example:

```
foo.s[0] = 1.0f; // Sets the 1st vector component of foo
pos.s[6] = 2; // Sets the 7th vector component of bar
```

In some cases vector components may also be accessed using the following notations. These notations are not guaranteed to be supported on all implementations, so their use should be accompanied by a check of the corresponding preprocessor symbol.

C.5.1 Named vector components notation

Vector data type components may be accessed using the `.sN`, `.sn` or `.xyzw` field naming convention, similar to how they are used within the OpenCL language. Use of the `.xyzw` field naming convention only allows accessing of the first 4 component fields. Support of these notations is identified by the `CL_HAS_NAMED_VECTOR_FIELDS` preprocessor symbol. For example:

```
#ifdef CL_HAS_NAMED_VECTOR_FIELDS
    cl_float4 foo;
    cl_int16 bar;
    foo.x = 1.0f; // Set first component
    foo.s0 = 1.0f; // Same as above
    bar.z = 3; // Set third component
    bar.se = 11; // Same as bar.s[0xe]
    bar.sD = 12; // Same as bar.s[0xd]
#endif
```

Unlike the OpenCL language type usage of named vector fields, only one component field may be accessed at a time. This restriction prevents the ability to swizzle or replicate components as is possible with the OpenCL language types. Attempting to access beyond the number of components for a type also results in a failure.

```
foo.xy // illegal - illegal field name combination
bar.s1234 // illegal - illegal field name combination
foo.s7 // illegal - no component s7
```

C.5.2 High/Low vector component notation

Vector data type components may be accessed using the `.hi` and `.lo` notation similar to that supported within the language types. Support of this notation is identified by the `CL_HAS_HI_LO_VECTOR_FIELDS` preprocessor symbol. For example:

```
#ifdef CL_HAS_HI_LO_VECTOR_FIELDS
    cl_float4 foo;
    cl_float2 new_hi = 2.0f, new_lo = 4.0f;
    foo.hi = new_hi;
    foo.lo = new_lo;
#endif
```


C.5.3 Native vector type notation

Certain native vector types are defined for providing a mapping of vector types to architecturally builtin vector types. Unlike the above described application vector types, these native types are supported on a limited basis depending on the supporting architecture and compiler.

These types are not unions, but rather convenience mappings to the underlying architectures' builtin vector types. The native types share the name of their application counterparts but are preceded by a double underscore "__".

For example, `__cl_float4` is the native builtin vector type equivalent of the `cl_float4` application vector type. The `__cl_float4` type may provide direct access to the architectural builtin `__m128` or vector float type, whereas the `cl_float4` is treated as a union.

In addition, the above described application data types may have native vector data type members for access convenience. The native components are accessed using the `.vN` sub-vector notation, where `N` is the number of elements in the sub-vector. In cases where the native type is a subset of a larger type (more components), the notation becomes an index based array of the sub-vector type.

Support of the native vector types is identified by a `__CL_TYPEN__` preprocessor symbol matching the native type name. For example:

```
#ifdef __CL_FLOAT4__ // Check for native cl_float4 type
    cl_float8 foo;
    __cl_float4 bar; // Use of native type
    bar = foo.v4[1]; // Access the second native float4
    vector
#endif
```

C.6 Implicit Conversions

Implicit conversions between application vector types are not supported.

C.7 Explicit Casts

Explicit casting of application vector types (`cl_typen`) is not supported. Explicit casting of native vector types (`__cl_typen`) is defined by the external compiler.

C.8 Other operators and functions

The behavior of standard operators and function on both application vector types (`cl_typen`) and native vector types (`__cl_typen`) is defined by the external compiler.

C.9 Application constant definitions

In addition to the above application type definitions, the following literal definitions are also available.

| | |
|--------------------------------|---|
| <code>CL_CHAR_BIT</code> | Bit width of a character |
| <code>CL_SCHAR_MAX</code> | Maximum value of a type <code>cl_char</code> |
| <code>CL_SCHAR_MIN</code> | Minimum value of a type <code>cl_char</code> |
| <code>CL_CHAR_MAX</code> | Maximum value of a type <code>cl_char</code> |
| <code>CL_CHAR_MIN</code> | Minimum value of a type <code>cl_char</code> |
| <code>CL_UCHAR_MAX</code> | Maximum value of a type <code>cl_uchar</code> |
| <code>CL_SHORT_MAX</code> | Maximum value of a type <code>cl_short</code> |
| <code>CL_SHORT_MIN</code> | Minimum value of a type <code>cl_short</code> |
| <code>CL_USHORT_MAX</code> | Maximum value of a type <code>cl_ushort</code> |
| <code>CL_INT_MAX</code> | Maximum value of a type <code>cl_int</code> |
| <code>CL_INT_MIN</code> | Minimum value of a type <code>cl_int</code> |
| <code>CL_UINT_MAX</code> | Maximum value of a type <code>cl_uint</code> |
| <code>CL_LONG_MAX</code> | Maximum value of a type <code>cl_long</code> |
| <code>CL_LONG_MIN</code> | Minimum value of a type <code>cl_long</code> |
| <code>CL_ULONG_MAX</code> | Maximum value of a type <code>cl_ulong</code> |
| | |
| <code>CL_FLT_DIAG</code> | Number of decimal digits of precision for the type <code>cl_float</code> |
| <code>CL_FLT_MANT_DIG</code> | Number of digits in the mantissa of type <code>cl_float</code> |
| <code>CL_FLT_MAX_10_EXP</code> | Maximum positive integer such that 10 raised to this power minus one can be represented as a normalized floating-point number of type <code>cl_float</code> |
| <code>CL_FLT_MAX_EXP</code> | Maximum exponent value of type <code>cl_float</code> |
| <code>CL_FLT_MIN_10_EXP</code> | Minimum negative integer such that 10 raised to this power minus one can be represented as a normalized floating-point number of type <code>cl_float</code> |
| <code>CL_FLT_MIN_EXP</code> | Minimum exponent value of type <code>cl_float</code> |
| <code>CL_FLT_RADIX</code> | Base value of type <code>cl_float</code> |
| <code>CL_FLT_MAX</code> | Maximum value of type <code>cl_float</code> |
| <code>CL_FLT_MIN</code> | Minimum value of type <code>cl_float</code> |
| <code>CL_FLT_EPSILON</code> | Minimum positive floating-point number of type |

| | |
|--------------------------------|--|
| | <code>cl_float</code> such that <code>1.0 + CL_FLT_EPSILON != 1</code> is true. |
| <code>CL_DBL_DIG</code> | Number of decimal digits of precision for the type <code>cl_double</code> |
| <code>CL_DBL_MANT_DIG</code> | Number of digits in the mantissa of type <code>cl_double</code> |
| <code>CL_DBL_MAX_10_EXP</code> | Maximum positive integer such that 10 raised to this power minus one can be represented as a normalized floating-point number of type <code>cl_double</code> |
| <code>CL_DBL_MAX_EXP</code> | Maximum exponent value of type <code>cl_double</code> |
| <code>CL_DBL_MIN_10_EXP</code> | Minimum negative integer such that 10 raised to this power minus one can be represented as a normalized floating-point number of type <code>cl_double</code> |
| <code>CL_DBL_MIN_EXP</code> | Minimum exponent value of type <code>cl_double</code> |
| <code>CL_DBL_RADIX</code> | Base value of type <code>cl_double</code> |
| <code>CL_DBL_MAX</code> | Maximum value of type <code>cl_double</code> |
| <code>CL_DBL_MIN</code> | Minimum value of type <code>cl_double</code> |
| <code>CL_DBL_EPSILON</code> | Minimum positive floating-point number of type <code>cl_double</code> such that <code>1.0 + CL_DBL_EPSILON != 1</code> is true. |
| | |
| <code>CL_NAN</code> | Macro expanding to a value representing NaN |
| <code>CL_HUGE_VALF</code> | Largest representative value of type <code>cl_float</code> |
| <code>CL_HUGE_VAL</code> | Largest representative value of type <code>cl_double</code> |
| <code>CL_MAXFLOAT</code> | Maximum value of type <code>cl_float</code> |
| <code>CL_INFINITY</code> | Macro expanding to a value representing infinity |

These literal definitions are defined in the `cl_platform.h` header.

Appendix D — OpenCL C++ Wrapper API

The OpenCL C++ wrapper API provides a C++ interface to the platform and runtime API. The C++ wrapper is built on top of the OpenCL 1.2 C API (platform and runtime) and is not a replacement. It is **required** that any implementation of the C++ wrapper API will make calls to the underlying C API and it is assumed that the C API is a compliant implementation of the OpenCL 1.2 specification.

Refer to the OpenCL C++ Wrapper API specification for details. The OpenCL C++ Wrapper API specification can be found at <http://www.khronos.org/registry/cl/>.

Appendix E — CL_MEM_COPY_OVERLAP

The following code describes how to determine if there is overlap between the source and destination rectangles specified to **clEnqueueCopyBufferRect** provided the source and destination buffers refer to the same buffer object.

Copyright (c) 2011 The Khronos Group Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and /or associated documentation files (the "Materials "), to deal in the Materials without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Materials, and to permit persons to whom the Materials are furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Materials.

THE MATERIALS ARE PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE MATERIALS OR THE USE OR OTHER DEALINGS IN THE MATERIALS.

```
bool
check_copy_overlap(size_t src_offset[3],
                  size_t dst_offset[3],
                  size_t region[3],
                  size_t row_pitch, size_t slice_pitch)
{
    const size_t src_min[] = {src_offset[0], src_offset[1], src_offset[2]};
    const size_t src_max[] = {src_offset[0] + region[0],
                              src_offset[1] + region[1],
                              src_offset[2] + region[2]};

    const size_t dst_min[] = {dst_offset[0], dst_offset[1], dst_offset[2]};
    const size_t dst_max[] = {dst_offset[0] + region[0],
                              dst_offset[1] + region[1],
                              dst_offset[2] + region[2]};

    // Check for overlap
    bool overlap = true;
    unsigned i;
    for (i=0; i != 3; ++i)
```

```

{
    overlap = overlap && (src_min[i] < dst_max[i])
                && (src_max[i] > dst_min[i]);
}

size_t dst_start = dst_offset[2] * slice_pitch +
                    dst_offset[1] * row_pitch + dst_offset[0];
size_t dst_end = dst_start + (region[2] * slice_pitch +
                               region[1] * row_pitch + region[0]);

size_t src_start = src_offset[2] * slice_pitch +
                    src_offset[1] * row_pitch + src_offset[0];
size_t src_end = src_start + (region[2] * slice_pitch +
                               region[1] * row_pitch + region[0]);

if (!overlap)
{
    size_t delta_src_x = (src_offset[0] + region[0] > row_pitch) ?
                        src_offset[0] + region[0] - row_pitch : 0;
    size_t delta_dst_x = (dst_offset[0] + region[0] > row_pitch) ?
                        dst_offset[0] + region[0] - row_pitch : 0;

    if ( (delta_src_x > 0 && delta_src_x > dst_offset[0]) ||
         (delta_dst_x > 0 && delta_dst_x > src_offset[0]) )
    {
        if ( (src_start <= dst_start && dst_start < src_end) ||
             (dst_start <= src_start && src_start < dst_end) )
            overlap = true;
    }
}

if (region[2] > 1)
{
    size_t src_height = slice_pitch / row_pitch;
    size_t dst_height = slice_pitch / row_pitch;

    size_t delta_src_y = (src_offset[1] + region[1] > src_height) ?
                        src_offset[1] + region[1] - src_height : 0;
    size_t delta_dst_y = (dst_offset[1] + region[1] > dst_height) ?
                        dst_offset[1] + region[1] - dst_height : 0;

    if ( (delta_src_y > 0 && delta_src_y > dst_offset[1]) ||
         (delta_dst_y > 0 && delta_dst_y > src_offset[1]) )
    {
        if ( (src_start <= dst_start && dst_start < src_end) ||
             (dst_start <= src_start && src_start < dst_end) )
            overlap = true;
    }
}
}

```

```
    }  
    return overlap;  
}
```

Appendix F – Changes

F.1 Summary of changes from OpenCL 1.0

The following features are added to the OpenCL 1.1 platform layer and runtime (*sections 4 and 5*):

- + Following queries to *table 4.3*
 - o CL_DEVICE_NATIVE_VECTOR_WIDTH_{CHAR | SHORT | INT | LONG | FLOAT | DOUBLE | HALF}
 - o CL_DEVICE_HOST_UNIFIED_MEMORY
 - o CL_DEVICE_OPENCL_C_VERSION
- + CL_CONTEXT_NUM_DEVICES to the list of queries specified to **clGetContextInfo**.
- + Optional image formats: CL_Rx, CL_RGx and CL_RGBx.
- + Support for sub-buffer objects – ability to create a buffer object that refers to a specific region in another buffer object using **clCreateSubBuffer**.
- + **clEnqueueReadBufferRect**, **clEnqueueWriteBufferRect** and **clEnqueueCopyBufferRect** APIs to read from, write to and copy a rectangular region of a buffer object respectively.
- + **clSetMemObjectDestructorCallback** API to allow a user to register a callback function that will be called when the memory object is deleted and its resources freed.
- + Options that control the OpenCL C version used when building a program executable. These are described in *section 5.6.4.5*.
- + CL_KERNEL_PREFERRED_WORK_GROUP_SIZE_MULTIPLE to the list of queries specified to **clGetKernelWorkGroupInfo**.
- + Support for user events. User events allow applications to enqueue commands that wait on a user event to finish before the command is executed by the device. Following new APIs are added - **clCreateUserEvent** and **clSetUserEventStatus**.
- + **clSetEventCallback** API to register a callback function for a specific command execution status.

The following modifications are made to the OpenCL 1.1 platform layer and runtime (*sections 4 and 5*):

- + Following queries in *table 4.3*

- CL_DEVICE_MAX_PARAMETER_SIZE from 256 to 1024 bytes
- CL_DEVICE_LOCAL_MEM_SIZE from 16 KB to 32 KB.

✚ The *global_work_offset* argument in **clEnqueueNDRangeKernel** can be a non-NULL value.

✚ All API calls except **clSetKernelArg** are thread-safe.

The following features are added to the OpenCL C programming language (*section 6*) in OpenCL 1.1:

✚ 3-component vector data types.

✚ New built-in functions

- **get_global_offset** work-item function defined in *section 6.12.1*.
- **minmag**, **maxmag** math functions defined in *section 6.12.2*.
- **clamp** integer function defined in *section 6.12.3*.
- (vector, scalar) variant of integer functions **min** and **max** in *section 6.12.3*.
- **async_work_group_strided_copy** defined in *section 6.12.10*.
- **vec_step**, **shuffle** and **shuffle2** defined in *section 6.12.12*.

✚ **cl_khr_byte_addressable_store** extension is a core feature.

✚ **cl_khr_global_int32_base_atomics**, **cl_khr_global_int32_extended_atomics**, **cl_khr_local_int32_base_atomics** and **cl_khr_local_int32_extended_atomics** extensions are core features. The built-in atomic function names are changed to use the **atomic_** prefix instead of **atom_**.

✚ Macros CL_VERSION_1_0 and CL_VERSION_1_1.

The following features in OpenCL 1.0 are deprecated in OpenCL 1.1:

- ✚ The **clSetCommandQueueProperty** API is no longer supported in OpenCL 1.1.
- ✚ The `__ROUNDING_MODE__` macro is no longer supported in OpenCL C 1.1.
- ✚ The `-cl-strict-aliasing` option that can be specified in *options* argument to **clBuildProgram** is no longer supported in OpenCL 1.1.

The following new extensions are added to *section 9* in OpenCL 1.1:

- ✚ **cl_khr_gl_event** – Creating a CL event object from a GL sync object.
- ✚ **cl_khr_d3d10_sharing** – Sharing memory objects with Direct3D 10.

The following modifications are made to the OpenCL ES Profile described in *section 10* in OpenCL 1.1:

✚ 64-bit integer support is optional.

F.2 Summary of changes from OpenCL 1.1

The following features are added to the OpenCL 1.2 platform layer and runtime (*sections 4 and 5*):

- ✚ Custom devices and built-in kernels are supported.
- ✚ Device partitioning that allows a device to be partitioned based on a number of partitioning schemes supported by the device.
- ✚ Extend *cl_mem_flags* to describe how the host accesses the data in a *cl_mem* object.
- ✚ **clEnqueueFillBuffer** and **clEnqueueFillImage** to support filling a buffer with a pattern or an image with a color.
- ✚ Add *CL_MAP_WRITE_INVALIDATE_REGION* to *cl_map_flags*. Appropriate clarification to the behavior of *CL_MAP_WRITE* has been added to the spec.
- ✚ New image types: 1D image, 1D image from a buffer object, 1D image array and 2D image arrays.
- ✚ **clCreateImage** to create an image object.
- ✚ **clEnqueueMigrateMemObjects** API that allows a developer to have explicit control over the location of memory objects or to migrate a memory object from one device to another.
- ✚ Support separate compilation and linking of programs.
- ✚ Additional queries to get the number of kernels and kernel names in a program have been added to **clGetProgramInfo**.
- ✚ Additional queries to get the compile and link status and options have been added to **clGetProgramBuildInfo**.
- ✚ **clGetKernelArgInfo** API that returns information about the arguments of a kernel.
- ✚ **clEnqueueMarkerWithWaitList** and **clEnqueueBarrierWithWaitList** APIs.

The following features are added to the OpenCL C programming language (*section 6*) in OpenCL 1.2:

- ✚ Double-precision is now an optional core feature instead of an extension.
- ✚ New built in image types: **image1d_t**, **image1d_array_t** and **image2d_array_t**.

- ✚ New built-in functions
 - Functions to read from and write to a 1D image, 1D and 2D image arrays described in *sections 6.12.14.2, 6.12.14.3 and 6.12.14.4.*
 - Sampler-less image read functions described in *section 6.12.14.3.*
 - **popcount** integer function described in *section 6.12.3.*
 - **printf** function described in *section 6.12.13.*

- ✚ Storage class specifiers `extern` and `static` as described in *section 6.8.*

- ✚ Macros `CL_VERSION_1_2` and `__OPENCL_C_VERSION__`.

The following APIs in OpenCL 1.1 are deprecated in OpenCL 1.2:

- ✚ **clEnqueueMarker**, **clEnqueueBarrier** and **clEnqueueWaitForEvents**
- ✚ **clCreateImage2D** and **clCreateImage3D**
- ✚ **clUnloadCompiler** and **clGetExtensionFunctionAddress**
- ✚ **clCreateFromGLTexture2D** and **clCreateFromGLTexture3D**

The following queries are deprecated in OpenCL 1.2:

- ✚ **CL_DEVICE_MIN_DATA_TYPE_ALIGN_SIZE** in *table 4.3* queried using **clGetDeviceInfo**.

Index - APIs

clBuildProgram, 138
clCompileProgram, 140
clCreateBuffer, 67
clCreateCommandQueue, 62
clCreateContext, 55
clCreateContextFromType, 57
clCreateImage, 91
clCreateKernel, 158
clCreateKernelsInProgram, 159
clCreateProgramWithBinary, 134
clCreateProgramWithBuiltInKernels, 136
clCreateProgramWithSource, 133
clCreateSampler, 129
clCreateSubBuffer, 70
clCreateSubDevices, 50
clCreateUserEvent, 180
clEnqueueBarrierWithWaitList, 189
clEnqueueCopyBuffer, 80
clEnqueueCopyBufferRect, 82
clEnqueueCopyBufferToImage, 111
clEnqueueCopyImage, 103
clEnqueueCopyImageToBuffer, 108
clEnqueueFillBuffer, 85
clEnqueueFillImage, 106
clEnqueueMapBuffer, 87
clEnqueueMapImage, 113
clEnqueueMarkerWithWaitList, 188
clEnqueueMigrateMemObjects, 124
clEnqueueNativeKernel, 176
clEnqueueNDRangeKernel, 171
clEnqueueReadBuffer, 73
clEnqueueReadBufferRect, 75
clEnqueueReadImage, 99
clEnqueueTask, 174
clEnqueueUnmapMemObject, 121
clEnqueueWriteBuffer, 73
clEnqueueWriteBufferRect, 76
clEnqueueWriteImage, 100
clFinish, 195
clFlush, 195
clGetCommandQueueInfo, 64
clGetContextInfo, 59
clGetDeviceIDs, 35
clGetDeviceInfo, 37
clGetEventInfo, 182
clGetEventProfilingInfo, 192
clGetImageInfo, 117
clGetKernelArgInfo, 167
clGetKernelInfo, 163
clGetKernelWorkGroupInfo, 165
clGetMemObjectInfo, 126
clGetPlatformIDs, 33
clGetPlatformInfo, 33
clGetProgramBuildInfo, 154
clGetProgramInfo, 151
clGetSamplerInfo, 131
clGetSupportedImageFormats, 97, 98
clLinkProgram, 143
clReleaseCommandQueue, 64
clReleaseContext, 58
clReleaseDevice, 53
clReleaseEvent, 187
clReleaseKernel, 160
clReleaseMemObject, 119
clReleaseProgram, 137
clReleaseSampler, 130
clRetainCommandQueue, 63
clRetainContext, 58
clRetainDevice, 53
clRetainEvent, 186
clRetainKernel, 160
clRetainMemObject, 119
clRetainProgram, 137
clRetainSampler, 130
clSetEventCallback, 185
clSetKernelArg, 161
clSetMemObjectDestructorCallback, 120
clSetUserEventStatus, 180
clUnloadPlatformCompiler, 150
clWaitForEvents, 181