

Reference Guide

**AMD Accelerated
Parallel Processing**
TECHNOLOGY



Graphics Core Next Architecture, Generation 3

August 2016

Revision 1.1

DISCLAIMER

The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

© 2016 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD Accelerated Parallel Processing, the AMD Accelerated Parallel Processing logo, and combinations thereof are trademarks of Advanced Micro Devices, Inc. OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos. Other names are for informational purposes only and may be trademarks of their respective owners.

AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.



Advanced Micro Devices, Inc.
One AMD Place
P.O. Box 3453
Sunnyvale, CA 94088-3453
www.amd.com

For AMD Accelerated Parallel Processing:

URL: developer.amd.com/appsdk
Developing: developer.amd.com/
Forum: developer.amd.com/openclforum

Contents

Contents

Preface

Chapter 1 Introduction

Chapter 2 Program Organization

2.1	The Compute Shader Program Type	2-1
2.2	Instruction Terminology	2-2
2.3	Data Sharing	2-3
2.3.1	Local Data Share (LDS)	2-4
2.3.2	Global Data Share (GDS)	2-5
2.4	Device Memory	2-5

Chapter 3 Kernel State

3.1	State Overview	3-1
3.2	Program Counter (PC)	3-2
3.3	EXECute Mask	3-2
3.4	Status Registers	3-2
3.5	Mode Register	3-4
3.6	GPRs and LDS	3-5
3.6.1	Out-of-Range Behavior	3-5
3.6.2	SGPR Allocation and Storage	3-6
3.6.3	SGPR Alignment	3-6
3.6.4	VGPR Allocation and Alignment	3-6
3.6.5	LDS Allocation and Clamping	3-6
3.7	M# Memory Descriptor	3-7
3.8	scc : Scalar Condition Code	3-7
3.9	Vector Compares: vcc and vccz	3-8
3.10	Trap and Exception Registers	3-8
3.11	Memory Violations	3-10

Chapter 4 Program Flow Control

4.1	Program Control	4-1
4.2	Branching	4-1
4.3	Work-Groups	4-2

4.4	Data Dependency Resolution	4-2
4.5	Manually Inserted Wait States (NOPs)	4-4
4.6	Arbitrary Divergent Control Flow	4-5
Chapter 5 Scalar ALU Operations		
5.1	SALU Instruction Formats	5-1
5.2	Scalar ALU Operands	5-2
5.3	Scalar Condition Code (SCC)	5-2
5.4	Integer Arithmetic Instructions	5-5
5.5	Conditional Instructions	5-5
5.6	Comparison Instructions	5-6
5.7	Bit-Wise Instructions	5-6
5.8	Special Instructions	5-8
Chapter 6 Vector ALU Operations		
6.1	Microcode Encodings	6-1
6.2	Operands	6-2
6.2.1	Instruction Inputs	6-2
6.2.2	Instruction Outputs	6-3
6.2.3	Out-of-Range GPRs	6-4
6.2.4	GPR Indexing	6-5
6.3	Instructions	6-5
6.4	Denormalized and Rounding Modes	6-7
6.5	ALU CLAMP Bit Usage	6-7
6.6	VGPR Indexing	6-7
6.6.1	Indexing Instructions	6-8
6.6.2	Special Cases	6-8
Chapter 7 Scalar Memory Operations		
7.1	Microcode Encoding	7-1
7.2	Operations	7-2
7.2.1	S_LOAD_DWORD	7-2
7.2.2	S_STORE_DWORD	7-2
7.2.3	S_BUFFER_LOAD_DWORD, S_BUFFER_STORE_DWORD	7-2
7.2.4	S_DCACHE_INV, S_DCACHE_WB	7-3
7.2.5	S_MEM_TIME	7-3
7.2.6	S_MEM_REALTIME	7-3
7.3	Dependency Checking	7-3
7.4	Alignment and Bounds Checking	7-3

Chapter 8 Vector Memory Operations

8.1	Vector Memory Buffer Instructions.....	8-1
8.1.1	Simplified Buffer Addressing	8-2
8.1.2	Buffer Instructions	8-2
8.1.3	VGPR Usage	8-4
8.1.4	Buffer Data	8-5
8.1.5	Buffer Addressing.....	8-6
	Range Checking8-8	
	Swizzled Buffer Addressing8-8	
	Proposed Uses Cases for Swizzled Addressing8-10	
8.1.6	Alignment	8-11
8.1.7	Buffer Resource	8-11
8.1.8	Memory Buffer Load to LDS	8-12
	Clamping Rules8-13	
8.1.9	GLC Bit Explained.....	8-13
8.2	Vector Memory (VM) Image Instructions.....	8-14
8.2.1	Image Instructions	8-15
8.2.2	Image Opcodes with No Sampler	8-16
8.2.3	Image Opcodes with Sampler.....	8-17
8.2.4	VGPR Usage	8-19
8.2.5	Image Resource.....	8-20
8.2.6	Sampler Resource.....	8-21
8.2.7	Data Formats	8-23
8.2.8	Vector Memory Instruction Data Dependencies.....	8-24

Chapter 9 Flat Memory Instructions

9.1	Flat Memory Instructions	9-1
9.2	Instructions.....	9-2
9.2.1	Ordering	9-3
9.2.2	Important Timing Consideration.....	9-3
9.3	Addressing.....	9-3
9.4	Memory Error Checking	9-3
9.5	Data.....	9-4
9.6	Scratch Space (Private).....	9-4

Chapter 10 Data Share Operations

10.1	Overview	10-1
10.2	Dataflow in Memory Hierarchy.....	10-2
10.3	LDS Access	10-3
10.3.1	LDS Direct Reads.....	10-3
10.3.2	LDS Parameter Reads	10-4
10.3.3	Data Share Indexed and Atomic Access.....	10-6

Chapter 11 Exporting Pixel Color and Vertex Shader Parameters

11.1	Microcode Encoding.....	11-1
11.2	Operations	11-2
11.2.1	Pixel Shader Exports	11-2
11.2.2	Vertex Shader Exports.....	11-2
11.3	Dependency Checking.....	11-3

Chapter 12 Instruction Set

12.1	SOP2 Instructions.....	12-1
12.2	SOPK Instructions	12-13
12.3	SOP1 Instructions.....	12-19
12.4	SOPC Instructions	12-32
12.5	SOPP Instructions.....	12-38
12.6	SMEM Instructions.....	12-47
12.7	VOP2 Instructions.....	12-58
12.8	VOP1 Instructions.....	12-83
12.9	VOPC Instructions	12-121
12.10	VOP3 3 in, 1 out Instructions (VOP3a).....	12-124
12.11	VOP3 Instructions (3 in, 2 out), (VOP3b).....	12-157
12.12	VINTRP Instructions	12-158
12.13	LDS/GDS Instructions.....	12-161
12.14	MUBUF Instructions.....	12-167
12.15	MTBUF Instructions.....	12-170
12.16	MIMG Instructions.....	12-173
12.17	EXP Instructions	12-177
12.18	FLAT Instructions.....	12-178

Chapter 13 Microcode Formats

13.1	Scalar ALU and Control Formats.....	13-3
13.2	Scalar Memory Instruction.....	13-15
13.3	Vector ALU instructions.....	13-18
13.4	Vector Parameter Interpolation Instruction.....	13-44
13.5	LDS/GDS Instruction.....	13-45
13.6	Vector Memory Buffer Instructions.....	13-50
13.7	Vector Memory Image Instruction.....	13-58
13.8	Export Instruction	13-62
13.9	FLAT Instruction.....	13-63

Figures

1.1 AMD Volcanic Islands Series Block Diagram1-1

1.2 Volcanic Islands Dataflow1-3

2.1 Shared Memory Hierarchy on the AMD Volcanic Islands Series of Stream Processors.....2-4

4.1 Example of Complex Control Flow Graph4-6

8.1 Buffer Address Components.....8-2

8.2 Address Calculation for a Linear Buffer8-8

8.3 Example of Buffer Swizzling8-10

8.4 Components of Addresses for LDS and Memory8-12

10.1 High-Level Memory Configuration10-1

10.2 Memory Hierarchy Dataflow10-2

10.3 LDS Layout with Parameters and Data Share.....10-5

Tables

2.1	Basic Instruction-Related Terms	2-2
2.2	Buffer, Texture, and Constant State	2-2
3.1	Readable and Writable Hardware States	3-1
3.2	Status Register Fields.....	3-3
3.3	Mode Register Fields.....	3-4
4.1	Control Instructions	4-1
4.2	Scalar ALU Instructions	4-1
4.3	VALU Instruction Set.....	4-4
5.1	Scalar Condition Code.....	5-4
5.2	Integer Arithmetic Instructions	5-5
5.3	Conditional Instructions.....	5-5
5.4	Comparison Instructions	5-6
5.5	Bit-Wise Instructions	5-6
5.6	Access Hardware Internal Register Instructions	5-8
5.7	Hardware Register Values	5-8
5.8	IB_STS	5-8
5.9	GPR_ALLOC.....	5-9
5.10	LDS_ALLOC.....	5-9
6.1	Instruction Operands.....	6-3
6.2	VALU Instruction Set.....	6-5
6.3	Compare Operations.....	6-6
6.4	MODE Register FP Bits.....	6-7
6.5	VGPR Indexing Instructions.....	6-8
7.1	SMRD Encoding Field Descriptions	7-1
8.1	Buffer Instructions	8-3
8.2	Microcode Formats	8-3
8.3	Address VGPRs.....	8-5
8.4	Buffer Instructions	8-5
8.5	BUFFER Instruction Fields for Addressing	8-6
8.6	V# Buffer Resource Constant Fields for Addressing	8-7
8.7	Address Components from GPRs	8-7
8.8	Address Components from GPRs	8-11
8.9	Buffer Resource Descriptor	8-11
8.10	Image Instructions.....	8-15
8.11	Instruction Fields.....	8-15
8.12	Image Opcodes with No Sampler	8-16
8.13	Image Opcodes with Sampler	8-17
8.14	Sample Instruction Suffix Key	8-18
8.15	Image Resource Definition	8-20
8.16	Sampler Resource Definition.....	8-21
8.17	Data and Image Formats.....	8-23
9.1	Flat Microcode Formats.....	9-2

10.1	Parameter Instruction Fields	10-5
10.2	LDS Instruction Fields	10-6
10.3	LDS Indexed Load/Store	10-7
11.1	EXP Encoding Field Descriptions	11-1
12.1	VOPC Instructions with 16 Compare Operations	12-122
12.2	VOPC Instructions with Eight Compare Operations	12-122
12.3	VOPC CLASS Instructions	12-123
12.4	Result of V_ADD_F64 Instruction	12-124
12.5	Result of LDEXP_F64 Instruction	12-139
12.6	Result of MUL_64 Instruction	12-150
12.7	DS Instructions for the Opcode Field	12-161
12.8	MUBUF Instructions for the Opcode Field	12-167
12.9	MTBUF Instructions for the Opcode Field	12-170
12.10	NFMT: Shader Num_Format	12-171
12.11	DFMT: Data_Format	12-171
12.12	MIMG Instructions for the Opcode Field	12-173
12.13	FLAT Instructions for the Opcode Field	12-178
13.1	Summary of Microcode Formats	13-1

Preface

About This Document

This document describes the environment, organization, and program state of AMD GCN Generation 3 devices. It details the instruction set and the microcode formats native to this family of processors that are accessible to programmers and compilers.

The document specifies the instructions (including the format of each type of instruction) and the relevant program state (including how the program state interacts with the instructions). Some instruction fields are mutually dependent; not all possible settings for all fields are legal. This document specifies the valid combinations.

The main purposes of this document are to:

1. Specify the language constructs and behavior, including the organization, of each type of instruction in both text syntax and binary format.
2. Provide a reference of instruction operation that compiler writers can use to maximize performance of the processor.

Audience

This document is intended for programmers writing application and system software, including operating systems, compilers, loaders, linkers, device drivers, and system utilities. It assumes that programmers are writing compute-intensive parallel applications (streaming applications) and assumes an understanding of requisite programming practices.

Organization

This document begins with an overview of the AMD GCN processors' hardware and programming environment ([Chapter 1](#)). [Chapter 2](#) describes the organization of GCN programs. [Chapter 3](#) describes the program state that is maintained. [Chapter 4](#) describes the program flow. [Chapter 5](#) describes the scalar ALU operations. [Chapter 6](#) describes the vector ALU operations. [Chapter 7](#) describes the scalar memory operations. [Chapter 8](#) describes the vector memory operations. [Chapter 9](#) provides information about the flat memory instructions. [Chapter 10](#) describes the data share operations. [Chapter 11](#) describes exporting the parameters of pixel color and vertex shaders. [Chapter 12](#) describes instruction details, first by the microcode format to which they belong, then in alphabetic order.

Finally, [Chapter 13](#) provides a detailed specification of each microcode format.

Conventions

The following conventions are used in this document.

mono-spaced font	A filename, file path, or code.
*	Any number of alphanumeric characters in the name of a code format, parameter, or instruction.
< >	Angle brackets denote streams.
[1,2)	A range that includes the left-most value (in this case, 1) but excludes the right-most value (in this case, 2).
[1,2]	A range that includes both the left-most and right-most values (in this case, 1 and 2).
{x y}	One of the multiple options listed. In this case, x or y.
0.0	A single-precision (32-bit) floating-point value.
1011b	A binary value, in this example a 4-bit value.
7:4	A bit range, from bit 7 to 4, inclusive. The high-order bit is shown first.
<i>italicized word or phrase</i>	The first use of a term or concept basic to the understanding of stream computing.

Related Documents

- *Intermediate Language (IL) Reference Manual*. Published by AMD.
- *AMD Accelerated Parallel Processing OpenCL Programming Guide*. Published by AMD.
- *The OpenCL Specification*. Published by Khronos Group. Aaftab Munshi, editor.
- *OpenGL Programming Guide*, at <http://www.khronos.org/registry/OpenGL/glew/>
- *Microsoft DirectX Reference Website*, at http://msdn.microsoft.com/archive/default.asp?url=/archive/en-us/directx9_c_Summer_04/directx/graphics/reference/reference.asp
- GPGPU: <http://www.gpgpu.org>

Differences Between GCN Generation 2 and 3 Devices

Important differences between Generation 2 and 3 GPUs

- Data Parallel ALU operations improve “Scan” and cross-lane operations.
- Scalar memory writes.
- In Generation 2, a kernel could read from a scalar data cache to retrieve constant data. In Generation 3, that cache now supports reads and writes.
- Compute kernel context switching.
- Compute kernels now can be context-switched on and off the GPU.

Summary of kernel instruction change from Generation 2 to 3

- Modified many of the microcode formats: VOP3A, VOP3B, LDS, GDS, MUBUF, MTBUF, MIMG, and EXP.
- SMRD microcode format is replaced with SMEM, now supporting reads and writes.
- VGPR Indexing for VALU instructions.
- New Instructions
 - Scalar Memory Writes.
 - S_CMP_EQ_U64, S_CMP_NE_U64.
 - 16-bit floating point VALU instructions.
 - “SDWA” – Sub Dword Addressing allows access to bytes and words of VGPRs in VALU instructions.
 - “DPP” – Data Parallel Processing allows VALU instructions to access data from neighboring lanes.
 - V_PERM_B32.
 - DS_PERMUTE_B32, DS_BPERMPUTE_B32.
- Removed Instructions
 - V_MAC_LEGACY_F32
 - V_CMPS* - now supported by V_CMP with the “clamp” bit set to 1.
 - V_MULLIT_F32.
 - V_{MIN, MAX, RCP, RSQ}_LEGACY_F32.
 - V_{LOG, RCP, RSQ}_CLAMP_F32.
 - V_{RCP, RSQ}_CLAMP_F64.
 - V_MUL_LO_I32 (it’s functionally identical to V_MUL_LO_U32).
 - All non-reverse shift instructions.
 - LDS and Memory atomics: MIN, MAX and CMPSWAP on F32 and F64 data.

- Removed Image Data Formats
 - snorm_lz (aka: snorm_ogl)
 - ubnorm
 - ubnorm_nz (aka: ubnorm_ogl)
 - ubint
 - ubscaled

Contact Information

For information concerning AMD Accelerated Parallel Processing developing, please see: developer.amd.com/.

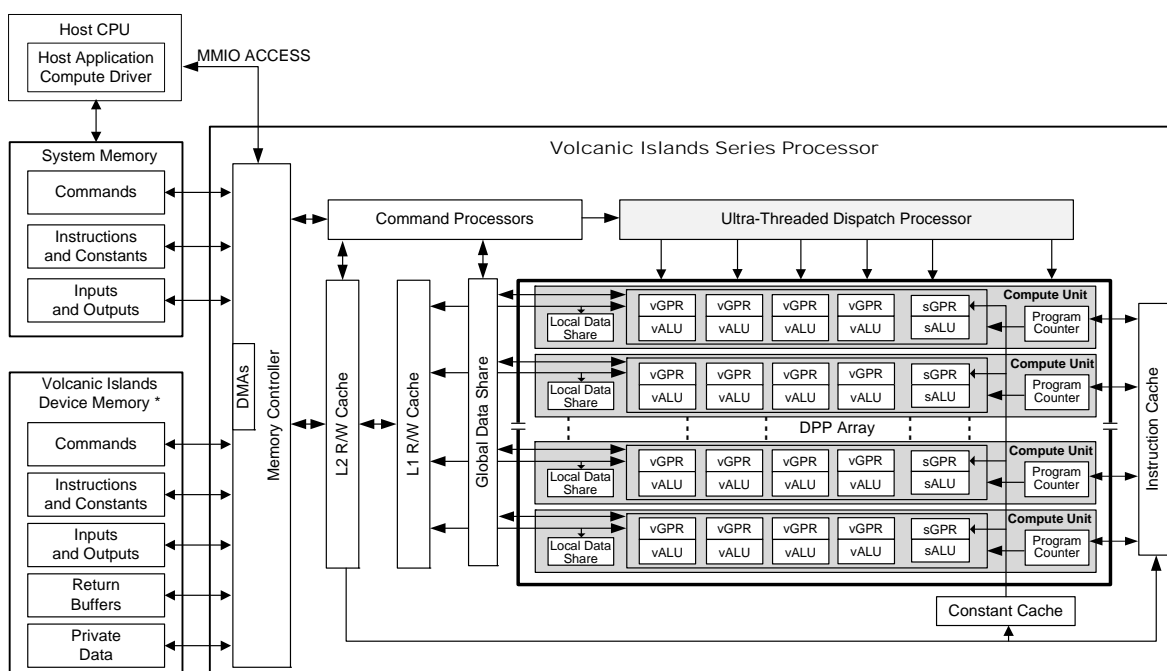
For information about developing with AMD Accelerated Parallel Processing, please see: developer.amd.com/appsdk.

We also have a growing community of AMD Accelerated Parallel Processing users. Come visit us at the AMD Accelerated Parallel Processing Developer Forum (developer.amd.com/openciforum) to find out what applications other users are trying on their AMD Accelerated Parallel Processing products.

Chapter 1 Introduction

AMD GCN processors implements a parallel microarchitecture that provides an excellent platform not only for computer graphics applications but also for general-purpose data parallel applications. Any data-intensive application that exhibits high bandwidth needs or significant computational requirements is a candidate for running on an AMD GCN processor.

Figure 1.1 shows a block diagram of the AMD GCN Generation 3 series processors.



*Discrete GPU – Physical Device Memory; APU – Region of system for GPU direct access

Figure 1.1 AMD GCN Generation 3 Series Block Diagram

It includes a data-parallel processor (DPP) array, a command processor, a memory controller, and other logic (not shown). The GCN command processor reads commands that the host has written to memory-mapped GCN registers in the system-memory address space. The command processor sends hardware-generated interrupts to the host when the command is completed. The GCN memory controller has direct access to all GCN device memory and the host-specified areas of system memory. To satisfy read and write requests, the

memory controller performs the functions of a direct-memory access (DMA) controller, including computing memory-address offsets based on the format of the requested data in memory.

A host application cannot write to the GCN device memory directly, but it can command the GCN device to copy programs and data between system memory and device memory. For the CPU to write to GPU memory, there are two ways:

- Request the GPU's DMA engine to write data by pointing to the location of the source data on CPU memory, then pointing at the offset in the GPU memory.
- Upload a kernel to run on the shaders that access the memory through the PCIe link, then process it and store it in the GPU memory.

In the GCN environment, a complete application includes two parts:

- a program running on the host processor, and
- programs, called *kernels*, running on the GCN processor.

The GCN programs are controlled by host commands, which

- set GCN internal base-address and other configuration registers,
- specify the data domain on which the GCN GPU is to operate,
- invalidate and flush caches on the GCN GPU, and
- cause the GCN GPU to begin execution of a program.

The GCN driver program runs on the host.

The DPP array is the heart of the GCN processor. The array is organized as a set of compute unit pipelines, each independent from the others, that operate in parallel on streams of floating-point or integer data. The compute unit pipelines can process data or, through the memory controller, transfer data to, or from, memory. Computation in a compute unit pipeline can be made conditional. Outputs written to memory can also be made conditional.

Host commands request a compute unit pipeline to execute a kernel by passing it:

- an identifier pair (x, y),
- a conditional value, and
- the location in memory of the kernel code.

When it receives a request, the compute unit pipeline loads instructions and data from memory, begins execution, and continues until the end of the kernel. As kernels are running, the GCN hardware automatically fetches instructions and data from memory into on-chip caches; GCN software plays no role in this. GCN software also can load data from off-chip memory into on-chip general-purpose registers (GPRs) and caches.

Conceptually, each compute unit pipeline maintains a separate interface to memory, consisting of index pairs and a field identifying the type of request (program instruction, floating-point constant, integer constant, boolean constant, input read, or output write). The index pairs for inputs, outputs, and constants are specified by the requesting GCN instructions from the hardware-maintained program state in the pipelines.

The AMD GCN devices can detect floating point exceptions and can generate interrupts. In particular, it detects IEEE floating-point exceptions in hardware; these can be recorded for post-execution analysis. The software interrupts shown in Figure 1.1 from the command processor to the host represent hardware-generated interrupts for signalling command-completion and related management functions.

Figure 1.2 shows the dataflow for a GCN application. For general-purpose applications, only one processing block performs all computation.



Figure 1.2 GCN Generation 3 Dataflow

The GCN processor hides memory latency by keeping track of potentially hundreds of work-items in different stages of execution, and by overlapping compute operations with memory-access operations.

Chapter 2

Program Organization

GCN programs consist of scalar instructions that operate on one value per wavefront, and vector instructions that operate on one value per thread. All program flow control is handled through scalar instructions. Vector ALU instructions operate on Vector General Purpose Registers (VGPRs) and can take up to three operands and produce both a value and mask result. The mask can be a condition code or a carry out. Scalar ALU instructions operate on Scalar GPRs and can take up to two inputs and produce a single integer or bit-mask output. Both vector and scalar instructions can modify the EXECute mask; this mask controls which threads are active and execute instructions at a given point in the kernel. Programs typically use instructions for fetching data through the texture cache for data loads.

2.1 The Compute Shader Program Type

The program type commonly run on the GCN GPU (see Figure 1.2, on page 1-3) is the Compute Shader (CS), which is a generic program (compute kernel) that uses an input work-item ID as an index to perform:

- gather reads on one or more sets of input data,
- arithmetic computation, and
- scatter writes to one or more set of output data to memory.

Compute shaders can write to multiple (up to eight) surfaces, which can be a mix of multiple render targets (MRTs), unordered access views (UAVs), and flat address space.

All program types accept the same instruction types, and all of the program types can run on any of the available DPP-array pipelines that support these programs; however, each kernel type has certain restrictions, which are described with that type.

2.2 Instruction Terminology

Table 2.1 summarizes some of the instruction-related terms used in this document. The instructions themselves are described in the remaining chapters. Details on each instruction are given in [Chapter 13, “Microcode Formats”](#).

Table 2.1 Basic Instruction-Related Terms

Term	Size (bits)	Description
Microcode format	32	One of several encoding formats for all instructions. They are described in Chapter 13, “Microcode Formats.”
Instruction	32 or 64	Every instruction is described with either 32 bits or 64 bits of microcode. <ul style="list-style-type: none"> • Vector Memory instructions are 64 bits. • Exports are 64 bits. • LDS and GDS are 64 bits. • Scalar ALU instructions are 32 bits but can have an additional 32 bits of literal constant data. • Vector ALU instructions can be 32 bits or 64 bits. The 32-bit versions can have an additional 32 bits of literal constant data.
Literal Constant	32	Literal constants specify a 32-bit constant value, either integer or float, in the instruction stream that supplies a value to a single 32-bit instruction.
Export	n/a	Copying, or the instruction to copy, from one or more VGPRs to one of the following output buffers: Pixel Color, Pixel Depth (Z), Vertex Parameter, or Vertex Position.
Fetch	n/a	Load data into VGPRs (vector fetch) or into SGPRs (scalar fetch), from the texture cache.
Quad	n/a	Four related pixels (for general-purpose programming: [x,y] data elements) in an aligned 2x2 space.
Fragment	n/a	A set of (x,y) data elements.
Pixel	n/a	A set of (x,y) data elements.
Thread	n/a	A work-item.

Table 2.2 summarizes the constant state data used by kernels for accessing memory buffer and image objects.

Table 2.2 Buffer, Texture, and Constant State

State	Access by GCN S/W	Access by Host S/W	Width (bits)	Description
Texture Samplers	R	W	128	A texture sampler describes how a texture map sample instruction (IMAGE) filters texel data and handles mip-maps. Texture samplers must first be loaded into four contiguous SGPRs prior to use by an IMAGE instruction.
Texture Resources	R	W	128 or 256	A texture resource describes the location, layout, and data type of a texture map in memory. A texture resource must be loaded into four or eight contiguous SGPRs prior to use by an IMAGE instruction.

Table 2.2 Buffer, Texture, and Constant State (Cont.)

State	Access by GCN S/W	Access by Host S/W	Width (bits)	Description
Constant Buffer	R	W	128	A specific usage of a buffer resource to describe an array of constant values that are provided by the application and loaded into memory prior to kernel invocation.
Border Color	No	W	128 (4 x 32 bits)	This is stored in the texture cache and referenced by texture samplers.

2.3 Data Sharing

The AMD GCN Stream processors can share data between different work-items. Data sharing can significantly boost performance. Figure 2.1 shows the memory hierarchy that is available to each work-item.

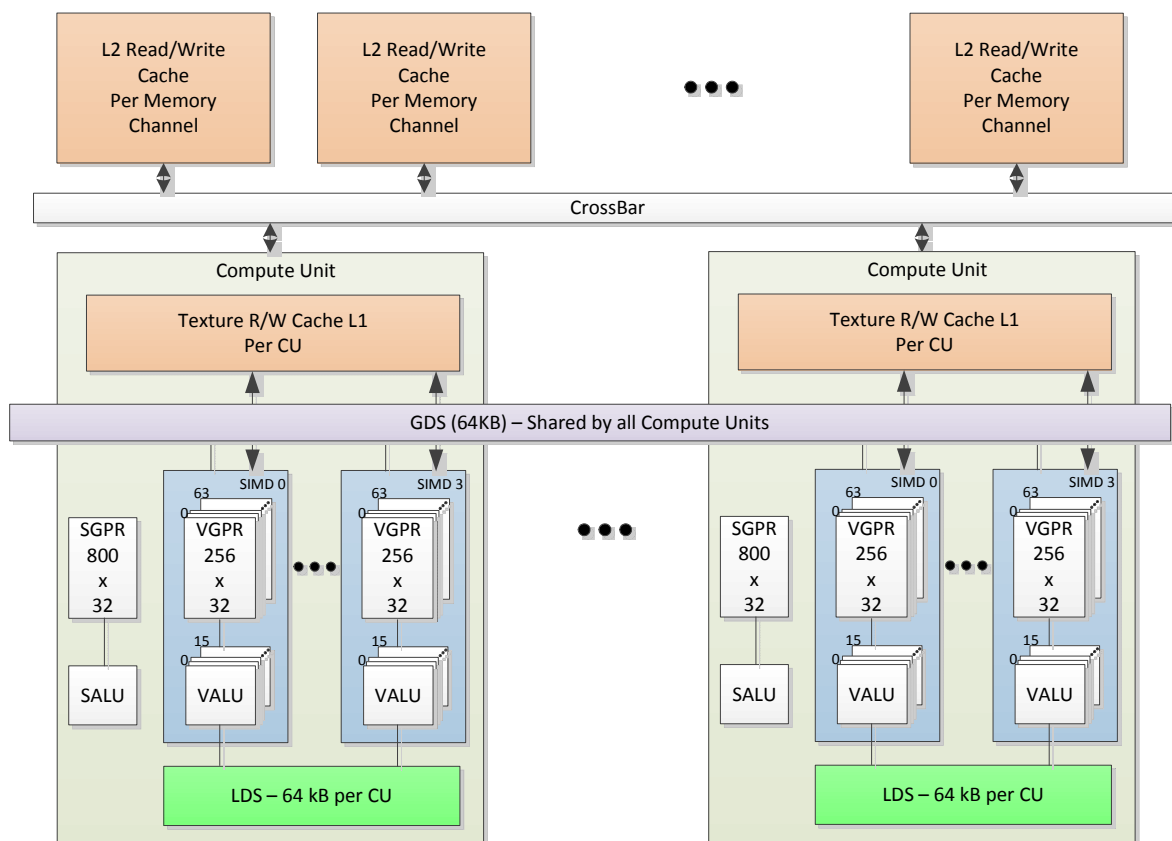


Figure 2.1 Shared Memory Hierarchy on the AMD GCN Generation 3 Series of Stream Processors

2.3.1 Local Data Share (LDS)

Each compute unit has a 32 kB memory space that enables low-latency communication between work-items within a work-group, or the work-items within a wavefront; this is the local data share (LDS). This memory is configured with

32 banks, each with 256 entries of 4 bytes. The AMD GCN processors use a 32 kB local data share (LDS) memory for each compute unit; this enables 128 kB of low-latency bandwidth to the processing elements. The AMD GCN devices have full access to any LDS location for any processor. The shared memory contains 32 integer atomic units to enable fast, unordered atomic operations. This memory can be used as a software cache for predictable re-use of data, a data exchange machine for the work-items of a work-group, or as a cooperative way to enable more efficient access to off-chip memory.

2.3.2 Global Data Share (GDS)

The AMD GCN devices use a 64 kB global data share (GDS) memory that can be used by wavefronts of a kernel on all compute units. This memory enables 128 bytes of low-latency bandwidth to all the processing elements. The GDS is configured with 32 banks, each with 512 entries of 4 bytes each. It provides full access to any location for any processor. The shared memory contains 32 integer atomic units to enable fast, unordered atomic operations. This memory can be used as a software cache to store important control data for compute kernels, reduction operations, or a small global shared surface. Data can be preloaded from memory prior to kernel launch and written to memory after kernel completion. The GDS block contains support logic for unordered append/consume and domain launch ordered append/consume operations to buffers in memory. These dedicated circuits enable fast compaction of data or the creation of complex data structures in memory.

2.4 Device Memory

The AMD GCN devices offer several methods for access to off-chip memory from the processing elements (PE) within each compute unit. On the primary read path, the device consists of multiple channels of L2 read-only cache that provides data to an L1 cache for each compute unit. Special cache-less load instructions can force data to be retrieved from device memory during an execution of a load clause. Load requests that overlap within the clause are cached with respect to each other. The output cache is formed by two levels of cache: the first for write-combining cache (collect scatter and store operations and combine them to provide good access patterns to memory); the second is a read/write cache with atomic units that lets each processing element complete unordered atomic accesses that return the initial value. Each processing element provides the destination address on which the atomic operation acts, the data to be used in the atomic operation, and a return address for the read/write atomic unit to store the pre-op value in memory. Each store or atomic operation can be set up to return an acknowledgement to the requesting PE upon write confirmation of the return value (pre-atomic op value at destination) being stored to device memory. This acknowledgement has two purposes:

- enabling a PE to recover the pre-op value from an atomic operation by performing a cache-less load from its return address after receipt of the write confirmation acknowledgement, and
- enabling the system to maintain a relaxed consistency model.

Each scatter write from a given PE to a given memory channel always maintains order. The acknowledgement enables one processing element to implement a fence to maintain serial consistency by ensuring all writes have been posted to memory prior to completing a subsequent write. In this manner, the system can maintain a relaxed consistency model between all parallel work-items operating on the system.

Chapter 3

Kernel State

This chapter describes the kernel states visible to the shader program.

3.1 State Overview

Table 3.1 describes all of the hardware states readable or writable by a shader program.

Table 3.1 Readable and Writable Hardware States

Abbrev.	Name	Size	Description
PC	Program Counter	40 bits	Points to the memory address of the next shader instruction to execute.
V0-V255	VGPR	32 bits	Vector general-purpose register.
S0-S101	SGPR	32 bits	Scalar general-purpose register.
LDS	Local Data Share	64 kB	Local data share is a scratch RAM with built-in arithmetic capabilities that allow data to be shared between threads in a workgroup.
EXEC	Execute Mask	64 bits	A bit mask with one bit per thread, which is applied to vector instructions and controls that threads execute and that ignore the instruction.
EXECZ	EXEC is zero	1 bit	A single bit flag indicating that the EXEC mask is all zeros.
VCC	Vector Condition Code	64 bits	A bit mask with one bit per thread; it holds the result of a vector compare operation.
VCCZ	VCC is zero	1 bit	A single bit-flag indicating that the VCC mask is all zeros.
SCC	Scalar Condition Code	1 bit	Result from a scalar ALU comparison instruction.
FLAT_SCRATCH	Flat scratch address	64 bits	The base address of scratch memory.
XNACK_MASK	Address translation failure.	64 bits	Bit mask of threads that have failed their address translation. Carrizo APU only.
STATUS	Status	32 bits	Read-only shader status bits.
MODE	Mode	32 bits	Writable shader mode bits.
M0	Memory Reg	32 bits	A temporary register that has various uses, including GPR indexing and bounds checking.
TRAPSTS	Trap Status	32 bits	Holds information about exceptions and pending traps.
TBA	Trap Base Address	64 bits	Holds the pointer to the current trap handler program.
TMA	Trap Memory Address	64 bits	Temporary register for shader operations. For example, can hold a pointer to memory used by the trap handler.

Table 3.1 Readable and Writable Hardware States (Cont.)

Abbrev.	Name	Size	Description
TTMP0- TTMP11	Trap Temporary SGPRs	32 bits	12 SGPRs available only to the Trap Handler for temporary storage.
VMCNT	Vector memory instruction count	4 bits	Counts the number of VMEM instructions issued but not yet completed.
EXPCNT	Export Count	3 bits	Counts the number of Export and GDS instructions issued but not yet completed. Also counts VMEM writes that have not yet sent their write-data to the TC.
LGKMCNT	LDS, GDS, Constant and Message count	4 bits	Counts the number of LDS, GDS, constant-fetch (scalar memory read), and message instructions issued but not yet completed.

3.2 Program Counter (PC)

The program counter (PC) is a byte address pointing to the next instruction to execute. When a wavefront is created, the PC is initialized to the first instruction in the program.

The PC interacts with three instructions: `S_GET_PC`, `S_SET_PC`, `S_SWAP_PC`. These transfer the PC to, and from, an even-aligned SGPR pair.

Branches jump to (`PC_of_the_instruction_after_the_branch + offset`). The shader program cannot directly read from, or write to, the PC. Branches, `GET_PC` and `SWAP_PC`, are PC-relative to the next instruction, not the current one. `S_TRAP` saves the PC of the `S_TRAP` instruction itself.

3.3 EXECute Mask

The Execute mask (64-bit) determines which threads in the vector are executed: 1 = execute, 0 = do not execute.

`EXEC` can be read from, and written to, through scalar instructions; it also can be written as a result of a vector-ALU compare. This mask affects vector-ALU, vector-memory, LDS, and export instructions. It does not affect scalar execution or branches.

A helper bit (`EXECZ`) can be used as a condition for branches to skip code when `EXEC` is zero.

Performance Note: unlike Evergreen, this hardware does no optimization when `EXEC = 0`. The shader hardware executes every instruction, wasting instruction issue bandwidth. Use `CBRANCH` or `VSKIP` to more rapidly skip over code when it is likely that the `EXEC` mask is zero.

3.4 Status Registers

Status register fields can be read, but not written to, by the shader. These bits are initialized at wavefront-creation time. Table 3.2 lists and briefly describes the status register fields.

Table 3.2 Status Register Fields

Field	Bit Position	Description
SCC	1	Scalar condition code. Used as a carry-out bit. For a comparison instruction, this bit indicates failure or success. For logical operations, this is 1 if the result was non-zero.
SPI_PRIO	2:1	Wavefront priority set by the shader processor interpolator (SPI) when the wavefront is created. See the <code>S_SETPRIO</code> instruction (page 12-45) for details. 0 is lowest, 3 is highest priority.
WAVE_PRIO	4:3	Wavefront priority set by the shader program. See the <code>S_SETPRIO</code> instruction (page 12-45) for details.
PRIV	5	Privileged mode. Can only be active when in the trap handler. Gives write access to the TTMP, TMA, and TBA registers.
TRAP_EN	6	Indicates that a trap handler is present. When set to zero, traps are never taken.
TTRACE_EN	7	Indicates whether thread trace is enabled for this wavefront. If zero, also ignore any shader-generated (instruction) thread-trace data.
EXPORT_RDY	8	This status bit indicates if export buffer space has been allocated. The shader stalls any export instruction until this bit becomes 1. It is set to 1 when export buffer space has been allocated. Before a Pixel or Vertex shader can export, the hardware checks the state of this bit. If the bit is 1, export can be issued. If the bit is zero, the wavefront sleeps until space becomes available in the export buffer. Then, this bit is set to 1, and the wavefront resumes.
EXECZ	9	Exec mask is zero.
VCCZ	10	Vector condition code is zero.
IN_TG	11	Wavefront is a member of a work-group of more than one wavefront.
IN_BARRIER	12	Wavefront is waiting at a barrier.
HALT	13	Wavefront is halted or scheduled to halt. HALT can be set by the host through wavefront-control messages, or by the shader. This bit is ignored while in the trap handler (<code>PRIV = 1</code>); it also is ignored if a host-initiated trap is received (request to enter the trap handler).
TRAP	14	Wavefront is flagged to enter the trap handler as soon as possible.
TTRACE_CU_EN	15	Enables/disables thread trace for this compute unit (CU). This bit allows more than one CU to be outputting USERDATA (shader initiated writes to the thread-trace buffer). Note that wavefront data is only traced from one CU per shader array. Wavefront user data (instruction based) can be output if this bit is zero.
VALID	16	Wavefront is active (has been created and not yet ended).
ECC_ERR	17	An ECC error has occurred.
SKIP_EXPORT	18	For Vertex Shaders only. 1 = this shader is never allocated export buffer space; all export instructions are ignored (treated as NOPs). Formerly called <code>VS_NO_ALLOC</code> . Used for stream-out of multiple streams (multiple passes over the same VS), and for DS running in the VS stage for wavefronts that produced no primitives.
PERF_EN	19	Performance counters are enabled for this wavefront.
COND_DBG_USER	20	Conditional debug indicator for user mode
COND_DBG_SYS	21	Conditional debug indicator for system mode.
ALLOW_REPLAY	22	Indicates that ATC replay is enabled.

Table 3.2 Status Register Fields (Cont.)

Field	Bit Position	Description
INST_ATC	23	Indicates the kernel instructions are located in ATC memory space. 0 = GPUVM.
DISPATCH_CACHE_CTRL	26:24	Indicates the cache policies for this dispatch. [24] = Vector L1 cache policy. [25] = L2 cache policy. [26] = Scalar data cache policy. The policies are: 0 = normal, 1 = force miss/evict for L1 and bypass for L2.
MUST_EXPORT	27	This wavefront is required to perform an export with Done=1 before terminating.

3.5 Mode Register

Mode register fields can be read from, and written to, by the shader through scalar instructions. Table 3.3 lists and briefly describes the mode register fields.

Table 3.3 Mode Register Fields

Field	Bit Position	Description
FP_ROUND	3:0	[1:0] Single precision round mode. [3:2] Double precision round mode. Round Modes: 0=nearest even, 1= +infinity, 2= -infinity, 3= toward zero.
FP_DENORM	7:4	[1:0] Single denormal mode. [3:2] Double denormal mode. Denorm modes: 0 = flush input and output denorms. 1 = allow input denorms, flush output denorms. 2 = flush input denorms, allow output denorms. 3 = allow input and output denorms.
DX10_CLAMP	8	Used by the vector ALU to force DX10-style treatment of NaNs: when set, clamp NaN to zero; otherwise, pass NaN through.
IEEE	9	Floating point opcodes that support exception flag gathering quiet and propagate signaling NaN inputs per IEEE 754-2008. Min_dx10 and max_dx10 become IEEE 754-2008 compliant due to signaling NaN propagation and quieting.
LOD_CLAMPED	10	Sticky bit indicating that one or more texture accesses had their LOD clamped.
DEBUG	11	Forces the wavefront to jump to the exception handler after each instruction is executed (but not after ENDPGM). Only works if TRAP_EN = 1.
EXCP_EN	18:12	Enable mask for exceptions. Enabled means if the exception occurs and TRAP_EN==1, a trap is taken. [12] : invalid. [15] : overflow. [13] : inputDenormal. [16] : underflow. [14] : float_div0. [17] : inexact. [18] : int_div0.
GPR_IDX_EN	27	GPR index enable.
VSKIP	28	0 = normal operation. 1 = skip (do not execute) any vector instructions: valu, vmem, export, lds, gds. "Skipping" instructions occurs at high-speed (10 wavefronts per clock cycle can skip one instruction). This is much faster than issuing and discarding instructions.
CSP	31:29	Conditional branch stack pointer. See Section 4.2 on page 4-1.

3.6 GPRs and LDS

This section describes how GPR and LDS space is allocated to a wavefront, as well as how out-of-range and misaligned accesses are handled.

3.6.1 Out-of-Range Behavior

When a source or destination is out of the legal range owned by a wavefront, the behavior is different from that resulting in the Northern Islands environment.

Out-of-range can occur through GPR-indexing or bad programming. It is illegal to index from one register type into another (for example: SGPRs into trap registers or inline constants). It is also illegal to index within inline constants.

The following describe the out-of-range behavior for various storage types.

- SGPRs
 - Source or destination out-of-range = (sgpr < 0 || (sgpr >= sgpr_size)).
 - Source out-of-range: returns the value of SGPR0 (not the value 0).
 - Destination out-of-range: instruction writes no SGPR result.
- VGPRs
 - Similar to SGPRs. It is illegal to index from SGPRs into VGPRs, or vice versa.
 - Out-of-range = (vgpr < 0 || (vgpr >= vgpr_size))
 - If a source VGPR is out of range, VGPR0 is used.
 - If a destination VGPR is out-of-range, the instruction is ignored (treated as an NOP).
- LDS
 - If the LDS-ADDRESS is out-of-range (addr < 0 or > (MIN(lds_size, m0))):
 - ◊ Writes out-of-range are discarded; it is undefined if SIZE is not a multiple of write-data-size.
 - ◊ Reads return the value zero.
 - If any source-VGPR is out-of-range, use the VGPR0 value is used.
 - If the dest-VGPR is out of range, nullify the instruction (issue with exec=0)
- Memory, LDS, and GDS: Reads and atomics with returns.
 - If any source VGPR or SGPR is out-of-range, the data value is undefined.
 - If any destination VGPR is out-of-range, the operation is nullified by issuing the instruction as if the EXEC mask were cleared to 0.
 - ◊ This out-of-range check must check all VGPRs that can be returned (for example: VDST to VDST+3 for a BUFFER_LOAD_DWORDx4).

- ◇ This check must also include the extra PRT (partially resident texture) VGPR and nullify the fetch if this VGPR is out-of-range, no matter whether the texture system actually returns this value or not.
- ◇ Atomic operations with out-of-range destination VGPRs are nullified: issued, but with exec mask of zero.

Instructions with multiple destinations (for example: `V_ADDC`): if any destination is out-of-range, no results are written.

3.6.2 SGPR Allocation and Storage

A wavefront can be allocated 8 to 102 SGPRs, in units of 16 GPRs (Dwords). These are logically viewed as SGPRs 0–101. The VCC is physically stored as part of the wavefront's SGPRs in the highest numbered two SGPRs (the source/destination VCC is an alias for those two SGPRs). When a trap handler is present, 16 additional SGPRs are reserved after VCC to hold the trap addresses, as well as saved-PC and trap-handler temps. These all are privileged (cannot be written to unless privilege is set). Note that if a wavefront allocates 16 SGPRs, 2 SGPRs are normally used as VCC, the remaining 14 are available to the shader. Shader hardware does not prevent use of all 16 SGPRs.

3.6.3 SGPR Alignment

Even-aligned SGPRs are required in the following cases.

- When 64-bit data is used. This is required for moves to/from 64-bit registers, including the PC.
- When scalar memory reads that the address-base comes from an SGPR-pair (either in SGPR).

Quad-alignment is required for the data-GPR when a scalar memory read returns four or more dwords.

When a 64-bit quantity is stored in SGPRs, the LSBs are in `SGPR[n]`, and the MSBs are in `SGPR[n+1]`.

3.6.4 VGPR Allocation and Alignment

VGPRs are allocated in groups of four Dwords. Operations using pairs of VGPRs (for example: double-floats) have no alignment restrictions. Physically, allocations of VGPRs can wrap around the VGPR memory pool.

3.6.5 LDS Allocation and Clamping

LDS is allocated per work-group or per-wavefront when work-groups are not in use. LDS space is allocated to a work-group or wavefront in contiguous blocks of 64 Dwords on 64-Dword alignment.

LDS allocations do not wrap around the LDS storage.

All accesses to LDS are restricted to the space allocated to that wavefront/work-group.

Clamping of LDS reads and writes is controlled by two size registers, which contain values for the size of the LDS space allocated by SPI to this wavefront or work-group, and a possibly smaller value specified in the LDS instruction (size is held in M0). The LDS operations use the smaller of these two sizes to determine how to clamp the read/write addresses.

3.7 M# Memory Descriptor

There is one 32-bit M# (M0) register per wavefront, which can be used for:

- Local Data Share (LDS)
 - Interpolation: holds { 1'b0, new_prim_mask[15:1], parameter_offset[15:0] } // in bytes
 - LDS direct-read offset and data type: { 13'b0, DataType[2:0], LDS_address[15:0] } // addr in bytes
 - LDS addressing for Memory/Vfetch → LDS: {16'h0, lds_offset[15:0]} // in bytes
 - Indexed LDS: provides SIZE in bytes { 15'h0, size[16:0] } // size in bytes
- Global Data Share (GDS)
 - { base[15:0] , size[15:0] } // base and size are in bytes
- Indirect GPR addressing for both vector and scalar instructions. M0 is an unsigned index.
- Send-message value. *EMIT/CUT* use M0 and EXEC as the send-message data.
- Flat: M0 provides the LDS SIZE in bytes (same as LDS-indexed case).

3.8 scc: Scalar Condition Code

Most scalar ALU instructions set the Scalar Condition Code (SCC) bit, indicating the result of the operation.

Compare operations: 1 = true

Arithmetic operations: 1 = carry out

Bit/logical operations: 1 = result was not zero

Move: does not alter SCC

The SCC can be used as the carry-in for extended-precision integer arithmetic, as well as the selector for conditional moves and branches.

3.9 Vector Compares: vcc and vccz

Vector ALU comparisons always set the Vector Condition Code (VCC) register (1=pass, 0=fail). Also, vector compares have the option of setting EXEC to the VCC value.

There is also a VCC summary bit (vccz) that is set to 1 when the VCC result is zero. This is useful for early-exit branch tests. VCC is also set for selected integer ALU operations (carry-out).

Vector compares have the option of writing the result to VCC (32-bit instruction encoding) or to any SGPR (64-bit instruction encoding). VCCZ is updated every time VCC is updated: vector compares and scalar writes to VCC.

The EXEC mask determines which threads execute an instruction. The VCC indicates which executing threads passed the conditional test, or which threads generated a carry-out from an integer add or subtract.

$$V_CMP_* \rightarrow VCC[n] = EXEC[n] \& (\text{test passed for thread}[n])$$

VCC is always fully written; there are no partial mask updates.

NOTE: VCC physically resides in the SGPR register file, so when an instruction sources VCC, that counts against the limit on the total number of SGPRs that can be sourced for a given instruction. VCC physically resides in the highest two user SGPRs.

Shader Hazard with VCC The user/compiler must prevent a scalar-ALU write to the SGPR holding VCC, immediately followed by a conditional branch using VCCZ. The hardware cannot detect this, and inserts the one required wait state (hardware *does* detect it when the SALU writes to VCC, it only fails to do this when the SALU instruction references the SGPRs that happen to hold VCC).

3.10 Trap and Exception Registers

Each type of exception can be enabled or disabled independently by setting, or clearing, bits in the TRAPSTS register's EXCP_EN field. This section describes the registers which control and report kernel exceptions.

All trap SGPRS (TBA, TMA, TTMP) are privileged for writes – they can be written only when in the trap handler (status.priv = 1). When not privileged, writes to these are ignored.

When a trap is taken (either user initiated, exception or host initiated), the shader hardware generates an S_TRAP instruction. This loads trap information into a pair of SGPRS:

$$\{TTMP1, TTMP0\} = \{3'h0, pc_rewind[3:0], HT[0], trapID[7:0], PC[47:0]\}.$$

“HT” is set to one for host initiated traps, and zero for user traps (s_trap) or exceptions. TRAP_ID is zero for exceptions, or the user/host trapID for those

traps. When the trap handler is entered, the PC of the faulting instruction is: (PC – PC_rewind*4).

STATUS . TRAP_EN – This bit indicates to the shader whether or not a trap handler is present. When one is not present, traps are never taken, no matter whether they’re floating point, user-, or host-initiated traps. When the trap handler is present, the wavefront uses an extra 16 SGPRs for trap processing. If trap_en == 0, all traps and exceptions are ignored, and s_trap is converted by hardware to NOP.

MODE . EXCP_EN[8:0] – Floating point exception enables. Defines which exceptions and events cause a trap.

Bit	Exception
0	invalid
1	Input Denormal
2	Divide by zero
3	overflow
4	underflow
5	inexact
6	integer divide by zero
7	address watch - TC (L1) has witnessed a thread access an ‘address of interest’
8	memory violation - a memory violation has occurred for this wave from L1 or LDS.

TRAP_STS Register –

Field	Bits	Description																				
EXCP	8:0	Status bits of which exceptions have occurred. These bits are sticky and accumulate results until the shader program clears them. These bits are accumulated regardless of the setting of EXCP_EN. These can be read or written without shader privilege. <table border="1"> <thead> <tr> <th>Bit</th> <th>Exception</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>invalid</td> </tr> <tr> <td>1</td> <td>Input Denormal</td> </tr> <tr> <td>2</td> <td>Divide by zero</td> </tr> <tr> <td>3</td> <td>overflow</td> </tr> <tr> <td>4</td> <td>underflow</td> </tr> <tr> <td>5</td> <td>inexact</td> </tr> <tr> <td>6</td> <td>integer divide by zero</td> </tr> <tr> <td>7</td> <td>address watch</td> </tr> <tr> <td>8</td> <td>memory violation</td> </tr> </tbody> </table>	Bit	Exception	0	invalid	1	Input Denormal	2	Divide by zero	3	overflow	4	underflow	5	inexact	6	integer divide by zero	7	address watch	8	memory violation
Bit	Exception																					
0	invalid																					
1	Input Denormal																					
2	Divide by zero																					
3	overflow																					
4	underflow																					
5	inexact																					
6	integer divide by zero																					
7	address watch																					
8	memory violation																					
SAVECTX	10	A bit set by the host command indicating that this wave must jump to its trap handler and save its context. This bit must be cleared by the trap handler using S_SETREG. Note – a shader can set this bit to 1 to cause a save-context trap, and (due to hardware latency) the shader can execute up to 2 additional instructions before taking the trap.																				

EXP_CYCLE	21:16	<p>When a float exception occurs, this tells the trap handler on which cycle the exception occurred on. 0-3 for normal float operations, 0-7 for double float add, and 0-15 for double float muladd or transcendentals. This register records the cycle number of the first occurrence of an enabled (unmasked) exception.</p> <table border="0"> <thead> <tr> <th>Field</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>EXCP_CYCLE[1:0]</td> <td>Phase: threads 0-15 are in phase 0, 48-63 in phase 3.</td> </tr> <tr> <td>EXCP_CYCLE[3:2]</td> <td>Multi-slot pass.</td> </tr> <tr> <td>EXCP_CYCLE[5:4]</td> <td>Hybrid pass – used for machines running at lower rates.</td> </tr> </tbody> </table>	Field	Meaning	EXCP_CYCLE[1:0]	Phase: threads 0-15 are in phase 0, 48-63 in phase 3.	EXCP_CYCLE[3:2]	Multi-slot pass.	EXCP_CYCLE[5:4]	Hybrid pass – used for machines running at lower rates.
Field	Meaning									
EXCP_CYCLE[1:0]	Phase: threads 0-15 are in phase 0, 48-63 in phase 3.									
EXCP_CYCLE[3:2]	Multi-slot pass.									
EXCP_CYCLE[5:4]	Hybrid pass – used for machines running at lower rates.									
DP_RATE	31:29	Determines how the shader interprets the TRAP_STS.cycle. Different Vector Shader Processors (VSP) process instructions at different rates.								

3.11 Memory Violations

A Memory Violation is reported from:

- LDS access out of range: $0 < \text{addr} < \text{lds_size}$. This can occur for indexed and direct access.
- LDS alignment error.
- Memory read/write/atomic out-of-range.
- Memory read/write/atomic alignment error.
- Flat access where the address is invalid (does not fall in any aperture).
- Write to a read-only surface.
- GDS alignment or address range error.
- GWS operation aborted (semaphore or barrier not executed).

Memory violations are not reported for instruction or scalar-data accesses.

Memory Buffer to LDS does NOT return a memory violation if the LDS address is out of range, but masks off EXEC bits of threads that would go out of range.

When a memory access is in violation, the appropriate memory (LDS or TC) returns MEM_VIOL to the wave. This is stored in the wave's TRAPSTS.mem_viol bit. This bit is sticky, so once set to 1, it remains at 1 until the user clears it.

There is a corresponding exception enable bit (EXCP_EN.mem_viol). If this bit is set when the memory returns with a violation, the wave jumps to the trap handler.

Memory violations are not precise. The violation is reported when the LDS or TC processes the address; during this time, the wave can have processed many more instructions. When a mem_viol is reported, the Program Counter saved is that of the next instruction to execute; it has no relationship the faulting instruction.

Chapter 4

Program Flow Control

All program flow control is programmed using scalar ALU instructions. This includes loops, branches, subroutine calls, and traps. The program uses SGPRs to store branch conditions and loop counters. Constants can be fetched from the scalar constant cache directly into SGPRs.

4.1 Program Control

The instructions in Table 4.1 control the priority and termination of a shader program, as well as provide support for trap handlers.

Table 4.1 Control Instructions

Instruction	Description
S_ENDPGM	Terminates the wavefront. It can appear anywhere in the kernel and can appear multiple times.
S_ENDPGM_SAVED	Terminates the wavefront due to context save. It can appear anywhere in the kernel and can appear multiple times.
S_NOP	Does nothing; it can be repeated in hardware up to eight times.
S_TRAP	Jumps to the trap handler.
S_RFE	Returns from the trap handler
S_SETPRIO	Modifies the priority of this wavefront: 0=lowest, 3 = highest.
S_SLEEP	Causes the wavefront to sleep for 64 – 960 clock cycles.
S_SENMSG	Sends a message (typically an interrupt) to the host CPU.

4.2 Branching

Branching is done using one of the following scalar ALU instructions.

Table 4.2 Scalar ALU Instructions

Instruction	Description
S_BRANCH	Unconditional branch.
S_CBRANCH_<test>	Conditional branch. Branch only if <test> is true. Tests are VCCZ, VCCNZ, EXECZ, EXECNZ, SCCZ, and SCCNZ.
S_CBRANCH_CDBGSYS	Conditional branch, taken if the COND_DBG_SYS status bit is set.
S_CBRANCH_CDBGUSER	Conditional branch, taken if the COND_DBG_USER status bit is set.

Table 4.2 Scalar ALU Instructions (Cont.)

Instruction	Description
S_CBRANCH_CDBGSYS_AND_USER	Conditional branch, taken only if both COND_DBG_SYS and COND_DBG_USER are set.
S_SETPC	Directly set the PC from an SGPR pair.
S_SWAPP	Swap the current PC with an address in an SGPR pair.
S_GETPC	Retrieve the current PC value (does not cause a branch).
S_CBRANCH_FORK and S_CBRANCH_JOIN	Conditional branch for complex branching.
S_SETVSKIP	Set a bit that causes all vector instructions to be ignored. Useful alternative to branching.

For conditional branches, the branch condition can be determined by either scalar or vector operations. A scalar compare operation sets the Scalar Condition Code (SCC), which then can be used as a conditional branch condition. Vector compare operations set the VCC mask, and VCCZ or VCCNZ then can be used to determine branching.

4.3 Work-Groups

Work-groups are collections of wavefronts running on the same compute unit which can synchronize and share data. Up to 16 wavefronts (1024 work-items) can be combined into a work-group. When multiple wavefronts are in a work-group, the S_BARRIER instruction can be used to force each wavefront to wait until all other wavefronts reach the same instruction; then, all wavefronts continue. Any wavefront can terminate early using S_ENDPGM, and the barrier is considered satisfied when the remaining live waves reach their barrier instruction.

4.4 Data Dependency Resolution

Shader hardware resolves most data dependencies, but a few cases must be explicitly handled by the shader program. In these cases, the program must insert S_WAITCNT instructions to ensure that previous operations have completed before continuing.

The shader has three counters that track the progress of issued instructions. S_WAITCNT waits for the values of these counters to be at, or below, specified values before continuing.

These allow the shader writer to schedule long-latency instructions, execute unrelated work, and specify when results of long-latency operations are needed.

Instructions of a given type return in order, but instructions of different types can complete out-of-order. For example, both GDS and LDS instructions use LGKM_cnt, but they can return out-of-order.

- **VM_CNT** – Vector memory count.
Determines when memory reads have returned data to VGPRs, or memory writes have completed.
 - Incremented every time a vector-memory read or write (MIMG, MUBUF, or MTBUF format) instruction is issued.
 - Decrementd for reads when the data has been written back to the VGPRs, and for writes when the data has been written to the L2 cache.

Ordering: Memory reads and writes return in the order they were issued, including mixing reads and writes.

- **LGKM_CNT** (LDS, GDS, (K)constant, (M)essage)
Determines when one of these low-latency instructions have completed.
 - Incremented by 1 for every LDS or GDS instruction issued, as well as by Dword-count for scalar-memory reads. For example, `s_memtime` counts the same as an `s_load_dwordx2`.
 - Decrementd by 1 for LDS/GDS reads or atomic-with-return when the data has been returned to VGPRs.
 - Incremented by 1 for each `S_SENDSMSG` issued. Decrementd by 1 when message is sent out.
 - Decrementd by 1 for LDS/GDS writes when the data has been written to LDS/GDS.
 - Decrementd by 1 for each Dword returned from the data-cache (SMEM).

Ordering

- ◊ Instructions of different types are returned out-of-order.
- ◊ Instructions of the same type are returned in the order they were issued, except scalar-memory-reads, which can return out-of-order (in which case only `S_WAITCNT 0` is the only legitimate value).

- **EXP_CNT** – VGPR-export count.
Determines when data has been read out of the VGPR and sent to GDS, at which time it is safe to overwrite the contents of that VGPR.
 - Incremented when an Export/GDS instruction is issued from the wavefront buffer.
 - Decrementd for exports/GDS when the last cycle of the export instruction is granted and executed (VGPRs read out).

Ordering

- ◊ Exports are kept in order only within each export type (color/null, position, parameter cache).

4.5 Manually Inserted Wait States (NOPs)

The hardware does not check for the following dependencies; they must be resolved by inserting NOPs or independent instructions.

Table 4.3 Required Software-Inserted Wait States

First Instruction	Second Instruction	Wait	Notes
S_SETREG <*>	S_GETREG <same reg>	2	
S_SETREG <*>	S_SETREG <same reg>	2	
SET_VSKIP	S_GETREG MODE	2	reads VSKIP from MODE
S_SETREG MODE.vskip	any vector op	2	requires 2 nops or non-vector instructions.
VALU which sets VCC or EXEC	VALU which uses EXECZ or VCCZ as a data source	5	
VALU writes SGPR/VCC (readlane, cmp, add/sub, div_scale)	V_{READ,WRITE}LANE using that SGPR/VCC as the lane select	4	
VALU writes VCC (including v_div_scale)	V_DIV_FMAS	4	
FLAT_STORE_X3 FLAT_STORE_X4 FLAT_ATOMIC_{F}CMPSWAP_X2 BUFFER_STORE_DWORD_X3 BUFFER_STORE_DWORD_X4 BUFFER_STORE_FORMAT_XYZ BUFFER_STORE_FORMAT_XYZW BUFFER_ATOMIC_{F}CMPSWAP_X2 IMAGE_STORE_* > 64 bits IMAGE_ATOMIC_{F}CMPSWAP > 64bits	Write VGPRs holding write-data from those instructions.	1	BUFFER_STORE_* operations which use an SGPR for 'offset' do not require any wait states. IMAGE_STORE_* and IMAGE_{F}CMPSWAP* ops with more than 2 DMASK bits set require this 1 wait state. Ops which use a 256-bit T# do not need a wait state.
VALU writes SGPR	VMEM reads that SGPR	5	Hardware assumes that there is no dependency here. If the VALU writes the SGPR that is used by a VMEM, the user must add 5 wait states.
SALU writes M0	GDS, S_SENDMSG or S_TTRACE_DATA	1	
VALU writes VGPR	VALU DPP reads that VGPR	2	

Table 4.3 Required Software-Inserted Wait States (Cont.)

First Instruction	Second Instruction	Wait	Notes
VALU writes EXEC	VALU DPP op	5	ALU does not forward EXEC to DPP.
Mixed use of VCC: alias vs SGPR# v_readlane, v_readfirstlane v_cmp v_add*_i/u v_sub*_i/u v_div_scale_* (writes vcc)	VALU which reads VCC as a constant (not as a carry-in which is 0 wait states).	1	VCC can be accessed by name or by the logical SGPR which holds VCC. The data dependency check logic does not understand that these are the same register and do not prevent races.
S_SETREG TRAPSTS	RFE, RFE_restore	1	

4.6 Arbitrary Divergent Control Flow

In the GCN architecture, conditional branches are handled in one of the following ways.

1. `S_CBRANCH`

This case is used for simple control flow, where the decision to take a branch is based on a previous compare operation. This is the most common method for conditional branching.

2. `S_CBRANCH_I/G_FORK` and `S_CBRANCH_JOIN`

This method, intended for more complex, irreducible control flow graphs, is described in the rest of this section. The performance of this method is lower than that for `S_CBRANCH` on simple flow control; use it only when necessary.

Conditional Branch (CBR) graphs are grouped into self-contained code blocks, denoted by `FORK` at the entrance point, and `JOIN` and the exit point (see Figure 4.1). The shader compiler must add these instructions into the code. This method uses a six-deep stack and requires three SGPRs for each fork/join block. Fork/Join blocks can be hierarchically nested to any depth (subject to SGPR requirements); they also can coexist with other conditional flow control or computed jumps.

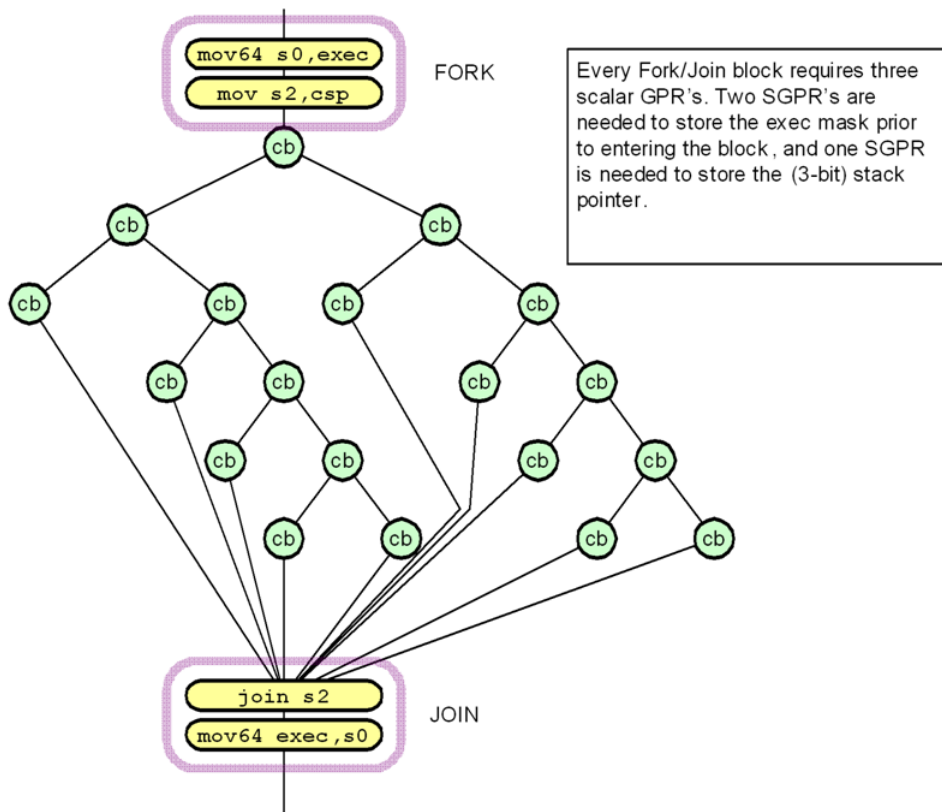


Figure 4.1 Example of Complex Control Flow Graph

The register requirements per wavefront are:

- CSP [2:0] - control stack pointer.
- Six stack entries of 128-bits each, stored in SGPRs: { exec[63:0], PC[47:2] }

This method compares how many of the 64 threads go down the PASS path instead of the FAIL path; then, it selects the path with the fewer number of threads first. This means at most 50% of the threads are active, and this limits the necessary stack depth to $\text{Log}_2 64 = 6$.

The following pseudo-code shows the details of CBRANCH Fork and Join operations.


```

S_CBRANCH_G_FORK arg0, arg1 // arg1 is an sgpr-pair which holds 64bit
                             (48bit) target address
S_CBRANCH_I_FORK arg0, #target_addr_offset[17:2] // target_addr_offset is a 16b signed
                                                 immediate offset
                                                 // "PC" in this pseudo-code is pointing to
                                                 the cbranch*_fork instruction

mask_pass = SGPR[arg0] & exec
mask_fail = ~SGPR[arg0] & exec
if (mask_pass == exec)
    I_FORK : PC += 4 + target_addr_offset
G_FORK: PC = SGPR[arg1]
    else if (mask_fail == exec)
        PC += 4
    else if (bitcount(mask_fail) < bitcount(mask_pass))
        exec = mask_fail
        I_FORK : SGPR[CSP*4] = { (pc + 4 + target_addr_offset), mask_pass }
G_FORK: SGPR[CSP*4] = { SGPR[arg1], mask_pass }
        CSP++
        PC += 4
    else
        exec = mask_pass
        SGPR[CSP*4] = { (pc+4), mask_fail }
        CSP++
        I_FORK : PC += 4 + target_addr_offset
G_FORK: PC = SGPR[arg1]

S_CBRANCH_JOIN arg0
if (CSP == SGPR[arg0]) // SGPR[arg0] holds the CSP value when the FORK started
    PC += 4 // this is the 2nd time to JOIN: continue with pgm
else
    CSP -- // this is the 1st time to JOIN: jump to other FORK path
    {PC, EXEC} = SGPR[CSP*4] // read 128-bits from 4 consecutive SGPRs

```

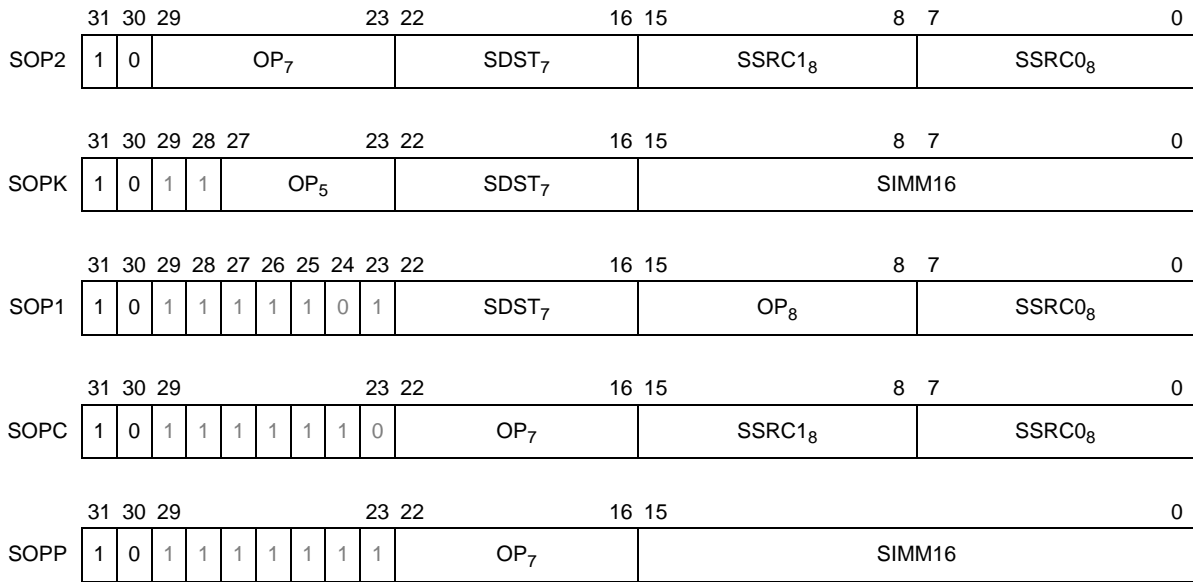

Chapter 5

Scalar ALU Operations

Scalar ALU (SALU) instructions operate on a single value per wavefront. These operations consist of 32-bit integer arithmetic and 32- or 64-bit bit-wise operations. The SALU also can perform operations directly on the Program Counter, allowing the program to create a call stack in SGPRs. Many operations also set the Scalar Condition Code bit (SCC) to indicate the result of a comparison, a carry-out, or whether the instruction result was zero.

5.1 SALU Instruction Formats

SALU instructions are encoded in one of 5 microcode formats, shown below:



Field	Description
OP	Opcode: instruction to be executed.
SDST	Destination SGPR.
SSRC0	First source operand.
SSRC1	Second source operand.
SIMM16	Signed immediate integer constant.

The lists of similar instructions sometimes use a condensed form using curly braces { } to express a list of possible names. For example, `S_AND_{B32, B64}` defines two legal instructions: `S_AND_B32` and `S_AND_B64`.

5.2 Scalar ALU Operands

Valid operands of SALU instructions are:

- SGPRs, including trap temporary SGPRs.
- Mode register.
- Status register (read-only).
- M0 register.
- TrapSts register.
- EXEC mask.
- VCC mask.
- SCC.
- PC.
- Inline constants: integers from -16 to 64, and a some floating point values.
- VCCZ, EXECZ, and SCC.
- Hardware registers.
- 32-bit literal constant.

The SALU cannot use VGPRs or LDS.

SALU instructions can use a 32-bit literal constant. This constant is part of the instruction stream and is available to all SALU microcode formats except SOPP and SOPK.

If any source SGPR is out-of-range, the value of SGPR0 is used instead.

If the destination SGPR is out-of-range, no SGPR is written with the result. However, SCC and possibly EXEC (if saveexec) are still written.

If an instruction uses 64-bit data in SGPRs, the SGPR pair must be aligned to an even boundary. For example, it is legal to use SGPRs 2 and 3 or 8 and 9 (but not 11 and 12) to represent 64-bit data.

5.3 Scalar Condition Code (SCC)

The scalar condition code (SCC) is written as a result of executing most SALU instructions.

The SCC is set by many instructions:

- Compare operations: 1 = true.
- Arithmetic operations: 1 = carry out.

- SCC = overflow for signed add and subtract operations. For add, overflow = both operands are of the same sign, and the MSB (sign bit) of the result is different than the sign of the operands. For subtract (A-B), overflow = A and B have opposite signs and the resulting sign is not the same as the sign of A.
- Bit/logical operations: 1 = result was not zero.

Table 5.1 Scalar Condition Code

		Code	Meaning	
Scalar Source (8 bits)	Scalar Dest (7 bits)	0 - 101	SGPR 0 to 103	Scalar GPRs.
		102	FLAT_SCR_LO	Holds the low Dword of the flat-scratch memory descriptor.
		103	FLAT_SCR_HI	Holds the high Dword of the flat-scratch memory descriptor.
		104	XNACK_MASK_LO	Holds the low Dword of the XNACK mask. Carrizo APU only.
		105	XNACK_MASK_HI	Holds the high Dword of the XNACK mask. Carrizo APU only.
		106	VCC_LO	vcc[31:0].
		107	VCC_HI	vcc[63:32].
		108	TBA_LO	Trap handler base address, [31:0].
		109	TBA_HI	Trap handler base address, [63:32].
		110	TMA_LO	Pointer to data in memory used by trap handler.
		111	TMA_HI	Pointer to data in memory used by trap handler.
		112-123	ttmp0 to ttmp11	Trap handler temps (privileged). {TTMP1, TTMP0} = {3'h0, PCRewind[3:0], HT[0], TrapID[7:0], PC[47:0]}
		124	M0	Temporary memory register.
		125	reserved	
		126	EXEC_LO	exec[31:0].
		127	EXEC_HI	exec[63:32].
		128	0	Immediate (constant value 0).
	129-192	int 1 to 64	Positive integer values.	
	193-208	int -1 to -16	Negative integer values.	
	209-239	reserved	unused	
	240	0.5	single or double floats	
	241	-0.5		
	242	1.0		
	243	-1.0		
	244	2.0		
	245	-2.0		
	246	4.0		
	247	-4.0		
	248-250	reserved	unused	
	251	VCCZ	{ zeros, VCCZ }	
	252	EXECZ	{ zeros, EXECZ }	
	253	SCC	{ zeros, SCC }	
254	reserved			
255	Literal constant	32-bit constant from instruction stream.		

5.4 Integer Arithmetic Instructions

This section describes the arithmetic operations supplied by the SALU.

Table 5.2 Integer Arithmetic Instructions

Instruction	Encoding	Sets SCC?	Operation
S_ADD_I32	SOP2	y	D = S1 + S2, SCC = overflow.
S_ADD_U32	SOP2	y	D = S1 + S2, SCC = carry out.
S_ADDC_U32	SOP2	y	D = S1 + S2 + SCC = overflow.
S_SUB_I32	SOP2	y	D = S1 - S2, SCC = overflow.
S_SUB_U32	SOP2	y	D = S1 - S2, SCC = carry out.
S_SUBB_U32	SOP2	y	D = S1 - S2 - SCC = carry out.
S_ABSDIFF_I32	SOP2	y	D = abs (s1 - s2), SCC = result not zero.
S_MIN_I32 S_MIN_U32	SOP2	y	D = (S1 < S2) ? S1 : S2. SCC = 1 if S1 was min.
S_MAX_I32 S_MAX_U32	SOP2	y	D = (S1 > S2) ? S1 : S2. SCC = 1 if S1 was max.
S_MUL_I32	SOP2	n	D = S1 * S2. Low 32 bits of result.
S_ADDK_I32	SOPK	y	D = D + simm16, SCC = overflow. Sign extended version of simm16.
S_MULK_I32	SOPK	n	D = D * simm16. Return low 32bits. Sign extended version of simm16.
S_ABS_I32	SOP1	y	D.i = abs (S1.i). SCC=result not zero.
S_SEXT_I32_I8	SOP1	n	D = { 24{S1[7]}, S1[7:0] }.
S_SEXT_I32_I16	SOP1	n	D = { 16{S1[15]}, S1[15:0] }.

5.5 Conditional Instructions

Conditional instructions use the SCC flag to determine whether to perform the operation, or (for CSELECT) which source operand to use.

Table 5.3 Conditional Instructions

Instruction	Encoding	Sets SCC?	Operation
S_CSELECT_{B32, B64}	SOP2	n	D = SCC ? S1 : S2.
S_CMOVK_I32	SOPK	n	if (SCC) D = signext(simm16).
S_CMOV_{B32, B64}	SOP1	n	if (SCC) D = S1, else NOP.

5.6 Comparison Instructions

These instructions compare two values and set the SCC to 1 if the comparison yielded a TRUE result.

Table 5.4 Comparison Instructions

Instruction	Encoding	Sets SCC?	Operation
S_CMP_EQ_U64, S_CMP_NE_U64	SOPC	y	Compare two 64-bit source values. SCC = S1 <cond> S2.
S_CMP_{EQ,NE,GT,GE,LE,LT}_{I32,U32}	SOPC	y	Compare two source values. SCC = S1 <cond> S2.
S_CMPK_{EQ,NE,GT,GE,LE,LT}_{I32,U32}	SOPK	y	Compare Dest SGPR to a constant. SCC = DST <cond> simm16. simm16 is zero-extended (U32) or sign-extended (I32).
S_BITCMP0_{B32,B64}	SOPC	y	Test for “is a bit zero”. SCC = !S1[S2].
S_BITCMP1_{B32,B64}	SOPC	y	Test for “is a bit one”. SCC = S1[S2].

5.7 Bit-Wise Instructions

Bit-wise instructions operate on 32- or 64-bit data without interpreting it as having a type. For bit-wise operations if noted in the table below, SCC is set if the result is nonzero.

Table 5.5 Bit-Wise Instructions

Instruction	Encoding	Sets SCC?	Operation
S_MOV_{B32,B64}	SOP1	n	D = S1
S_MOVK_I32	SOPK	n	D = signext(simm16)
{S_AND,S_OR,S_XOR}_{B32,B64}	SOP2	y	D = S1 & S2, S1 OR S2, S1 XOR S2
{S_ANDN2,S_ORN2}_{B32,B64}	SOP2	y	D = S1 & ~S2, S1 OR ~S2, S1 XOR ~S2,
{S_NAND,S_NOR,S_XNOR}_{B32,B64}	SOP2	y	D = ~(S1 & S2), ~(S1 OR S2), ~(S1 XOR S2)
S_LSHL_{B32,B64}	SOP2	y	D = S1 << S2[4:0], [5:0] for B64.
S_LSHR_{B32,B64}	SOP2	y	D = S1 >> S2[4:0], [5:0] for B64.
S_ASHR_{I32,I64}	SOP2	y	D = sext(S1 >> S2[4:0]) ([5:0] for I64).
S_BFM_{B32,B64}	SOP2	n	Bit field mask. D = ((1 << S1[4:0]) - 1) << S2[4:0].
S_BFE_U32, S_BFE_U64 S_BFE_I32, S_BFE_I64 (signed/unsigned)	SOP2	n	Bit Field Extract, then sign-extend result for I32/64 instructions. S1 = data, S2[5:0] = offset, S2[22:16] = width.
S_NOT_{B32,B64}	SOP1	y	D = ~S1.

Table 5.5 Bit-Wise Instructions (Cont.)

Instruction	Encoding	Sets SCC?	Operation
S_WQM_{B32,B64}	SOP1	y	D = wholeQuadMode(S1). If any bit in a group of four is set to 1, set the resulting group of four bits all to 1.
S_QUADMASK_{B32,B64}	SOP1	y	D[0] = OR(S1[3:0]), D[1]=OR(S1[7:4]), etc.
S_BREV_{B32,B64}	SOP1	n	D = S1[0:31] are reverse bits.
S_BCNT0_I32_{B32,B64}	SOP1	y	D = CountZeroBits(S1).
S_BCNT1_I32_{B32,B64}	SOP1	y	D = CountOneBits(S1).
S_FF0_I32_{B32,B64}	SOP1	n	D = Bit position of first zero in S1 starting from LSB. -1 if not found.
S_FF1_I32_{B32,B64}	SOP1	n	D = Bit position of first one in S1 starting from LSB. -1 if not found.
S_FLBIT_I32_{B32,B64}	SOP1	n	Find last bit. D = the number of zeros before the first one starting from the MSB. Returns -1 if none.
S_FLBIT_I32 S_FLBIT_I32_I64	SOP1	n	Count how many bits in a row (from MSB to LSB) are the same as the sign bit. Return -1 if the input is zero or all 1's (-1). 32-bit pseudo-code: <pre> if (S0 == 0 S0 == -1) D = -1 else D = 0 for (I = 31 .. 0) if (S0[I] == S0[31]) D++ else break </pre> This opcode behaves the same as V_FFBH_I32.
S_BITSET0_{B32,B64}	SOP1	n	D[S1[4:0], [5:0] for B64] = 0
S_BITSET1_{B32,B64}	SOP1	n	D[S1[4:0], [5:0] for B64] = 1
S_{and,or,xor,andn2,orn2,nand,nor,xnor} _SAVEEXEC_B64	SOP1	y	Save the EXEC mask, then apply a bit-wise operation to it. D = EXEC EXEC = S1 <op> EXEC SCC = (exec != 0)
S_MOVRELS_{B32,B64} S_MOVRELD_{B32,B64}	SOP1	n	Move a value into an SGPR relative to the value in M0. MOVERELS: D = SGPR[S1+M0] MOVERELD: SGPR[D+M0] = S1 Index must be even for 64. M0 is an unsigned index.

5.8 Special Instructions

These instructions access hardware internal registers.

Table 5.6 Access Hardware Internal Register Instructions

Instruction	Encoding	Sets SCC?	Operation
S_GETREG_B32	SOPK*	n	Read a hardware register into the LSBs of D.
S_SETREG_B32	SOPK*	n	Write the LSBs of D into a hardware register. (Note that D is a source SGPR.) Must add an S_NOP between two consecutive S_SETREG to the same register.
S_SETREG_IMM32_B32	SOPK*	n	S_SETREG where 32-bit data comes from a literal constant (so this is a 64-bit instruction format).

The hardware register is specified in the DEST field of the instruction, using the values in Table 5.7. Some bits of the DEST specify which register to read/write, but additional bits specify which bits in the special register to read/write:

SIMM16 = {size[4:0], offset[4:0], hwRegId[5:0]}; offset is 0..31, size is 1..32.

Table 5.7 Hardware Register Values

Code	Register	Description
0	reserved	
1	MODE	R/W.
2	STATUS	Read only.
3	TRAPSTS	R/W.
4	HW_ID	Read only. Debug only.
5	GPR_ALLOC	Read only. {sgpr_size, sgpr_base, vgpr_size, vgpr_base }
6	LDS_ALLOC	Read only. {lds_size, lds_base}
7	IB_STS	Read only. {valu_cnt, lgkm_cnt, exp_cnt, vm_cnt}

The following tables describe some of the registers in Table 5.7.

Table 5.8 IB_STS

Field	Bits	Description
VM_CNT	3:0	Number of VMEM instructions issued but not yet returned.

Table 5.8 IB_STS

Field	Bits	Description
EXP_CNT	6:4	Number of Exports issued but have not yet read their data from VGPRs.
LGKM_CNT	11:8	LDS, GDS, Constant-memory and Message instructions issued-but-not-completed count.
VALU_CNT	14:12	Number of VALU instructions outstanding for this wavefront.

Table 5.9 GPR_ALLOC

Field	Bits	Description
VGPR_BASE	5:0	Physical address of first VGPR assigned to this wavefront, as [7:2]
VGPR_SIZE	13:8	Number of VGPRs assigned to this wavefront, as [7:2]. 0=4 VGPRs, 1=8 VGPRs, etc.
SGPR_BASE	21:16	Physical address of first SGPR assigned to this wavefront, as [7:3].
SGPR_SIZE	27:24	Number of SGPRs assigned to this wave, as [7:3]. 0=8 SGPRs, 1=16 SGPRs, etc.

Table 5.10 LDS_ALLOC

Field	Bits	Description
LDS_BASE	7:0	Physical address of first LDS location assigned to this wavefront, in units of 64 Dwords.
LDS_SIZE	20:12	Amount of LDS space assigned to this wavefront, in units of 64 Dwords.

Chapter 6

Vector ALU Operations

Vector ALU instructions (VALU) perform an arithmetic or logical operation on data for each of 64 threads and write results back to VGPRs, SGPRs or the EXEC mask.

Parameter interpolation is a mixed VALU and LDS instruction, and is described in the Data Share chapter.


6.1 Microcode Encodings

Most VALU instructions are available in two encodings: VOP3 which uses 64-bits of instruction and has the full range of capabilities, and one of three 32-bit encodings that offer a restricted set of capabilities. A few instructions are only available in the VOP3 encoding. The only instructions that cannot use the VOP3 format are the parameter interpolation instructions.


When an instruction is available in two microcode formats, it is up to the user to decide which to use. It is recommended to use the 32-bit encoding whenever possible.

The microcode encodings are shown below.

VOP2 is for instructions with two inputs and a single vector destination. Instructions that have a carry-out implicitly write the carry-out to the VCC register.



VOP1 is for instructions with no inputs or a single input and one destination.



VOPC is for comparison instructions.



VINTRP is for parameter interpolation instructions.

1	1	0	1	0	1	VDST	OP	ATTR	ATTR CHAN	VSRC (I, J)	+0
---	---	---	---	---	---	------	----	------	--------------	-------------	----

VOP3 is for instructions with up to three inputs, input modifiers (negate and absolute value), and output modifiers. There are two forms of VOP3: one which uses a scalar destination field (used only for div_scale, integer add and subtract); this is designated VOP3b. All other instructions use the common form, designated VOP3a.

VOP3a:

NEG	OMOD	SRC2	SRC1	SRC0	+4						
1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0

VOP3b:

NEG	OMOD	SRC2	SRC1	SRC0	+4					
1	1	0	1	0	0	OP	CL MP	SDST	VDST	+0

Any of the 32-bit microcode formats can use a 32-bit literal constant, but not VOP3.

6.2 Operands

All VALU instructions take at least one input operand (except `V_NOP` and `V_CLREXCP`). The data-size of the operands is explicitly defined in the name of the instruction. For example, `V_MAD_F32` operates on 32-bit floating point data.

6.2.1 Instruction Inputs

VALU instructions can use any of the following sources for input, subject to restrictions listed below:

- VGPRs.
- SGPRs.
- Inline constants – a constant selected by a specific VSRC value (see Table 6.1).
- Literal constant – a 32-bit value in the instruction stream. When a literal constant is used with a 64bit instruction, the literal is expanded to 64 bits by:

padding the LSBs with zeros for floats, padding the MSBs with zeros for unsigned ints, and by sign-extending signed ints.

- LDS direct data read.
- M0.
- EXEC mask.

Limitations

- At most one SGPR can be read per instruction, but the value can be used for more than one operand.
- At most one literal constant can be used, and only when an SGPR or M0 is not used as a source.
- Only SRC0 can use LDS_DIRECT (see [Chapter 10, “Data Share Operations”](#)).

Instructions using the VOP3 form and also using floating-point inputs have the option of applying absolute value (ABS field) or negate (NEG field) to any of the input operands.

For both integer and floating point values: when a 32-bit literal constant is in a 64-bit instruction, the 32 bit constant is extended to 64 bits by padding the least significant 32 bits with zeroes. $1/(2*PI)$ is special for 64-bit extension: the mantissa bits from the 32-bit version are copied and the rest of the least significant mantissa bits are filled in with zero.

6.2.2 Instruction Outputs

VALU instructions typically write their results to VGPRs specified in the VDST field of the microcode word. A thread only writes a result if the associated bit in the EXEC mask is set to 1.

All V_CMPX instructions write the result of their comparison (one bit per thread) to both an SGPR (or VCC) and the EXEC mask.

Instructions producing a carry-out (integer add and subtract) write their result to VCC when used in the VOP2 form, and to an arbitrary SGPR-pair when used in the VOP3 form.

When the VOP3 form is used, instructions with a floating-point result can apply an output modifier (OMOD field) that multiplies the result by: 0.5, 1.0, 2.0 or 4.0. Optionally, the result can be clamped (CLAMP field) to the range [-1.0, +1.0], as indicated in Table 6.1.

In Table 6.1, all codes can be used when the vector source is nine bits; codes 0 to 255 can be the scalar source if it is eight bits; codes 0 to 127 can be the scalar source if it is seven bits; and codes 256 to 511 can be the vector source or destination.

Table 6.1 Instruction Operands

Field	Bit Position	Description
0 – 101	SGPR 0 .. 103	
102	FLATSCR_LO	Flat Scratch[31:0].
103	FLATSCR_HI	Flat Scratch[63:32].
104	XNACK_MASK_LO	Carrizo APU only.
105	XNACK_MASK_HI	Carrizo APU only.
106	VCC_LO	vcc[31:0].
107	VCC_HI	vcc[63:32].
108	TBA_LO	Trap handler base address, [31:0].
109	TBA_HI	Trap handler base address, [63:32].
110	TMA_LO	Pointer to data in memory used by trap handler.
111	TMA_HI	Pointer to data in memory used by trap handler.
112-123	ttmp0..ttmp11	Trap handler temps (privileged). {ttmp1,ttmp0} = PC_save{hi,lo}
124	M0	
125	reserved	
126	EXEC_LO	exec[31:0].
127	EXEC_HI	exec[63:32].
128	0	
129-192	int 1.. 64	Integer inline constants.
193-208	int -1 .. -16	
209-239	reserved	Unused.
240	0.5	Single, double, or half-precision inline floats. 1/(2*PI) is 0.15915494. The exact value used is: half: 0x3118 single: 0x3e22f983 double: 0x3fc45f306dc9c882
241	-0.5	
242	1.0	
243	-1.0	
244	2.0	
245	-2.0	
246	4.0	
247	-4.0	
248	1/(2*PI)	
249-250	reserved	Unused.
251	VCCZ	{ zeros, VCCZ }
252	EXECZ	{ zeros, EXECZ }
253	SCC	{ zeros, SCC }
254	LDS direct	Use LDS direct read to supply 32-bit value <i>Vector-alu instructions only.</i>
255	Literal constant	32-bit constant from instruction stream.
256 – 511	VGPR 0 .. 255	

6.2.3 Out-of-Range GPRs

When a source VGPR is out-of-range, the instruction uses as input the value from VGPR0.

When the destination GPR is out-of-range, the instruction executes but does not write the results.

6.3 Instructions

Table 6.2 lists the complete VALU instruction set by microcode encoding.

Table 6.2 VALU Instruction Set

VOP3	VOP3 – 1-2 operand opcodes	VOP2	VOP1
V_MAD_LEGACY_F32	V_ADD_F64	V_ADD_{ F16,F32, U16,U32}	V_NOP
V_MAD_{ F16,I16,U16,F32}	V_MUL_F64	V_SUB_{ F16,F32,U16, U32}	V_MOV_B32
V_MAD_I32_I24	V_MIN_F64	V_SUBREV_{ F16,F32, U16,U32}	V_READFIRSTLANE_B32
V_MAD_U32_U24	V_MAX_F64	V_ADDC_U32	V_CVT_F32_{I32,U32,F16,F64}
V_CUBEID_F32	V_LDEXP_F64	V_SUBB_U32	V_CVT_{I32,U32,F16, F64}_F32
V_CUBESC_F32	V_MUL_LO_U32	V_SUBBREV_U32	V_CVT_{I32,U32}_F64
V_CUBETC_F32	V_MUL_HI_{I32,U32}	V_MUL_LEGACY_F32	V_CVT_F64_{I32,U32}
V_CUBEMA_F32	V_LSHLREV_B64	V_MUL_{F16, F32}	V_CVT_F32_UBYTE{0,1,2,3}
V_BFE_{U32 , I32 }	V_LSHRREV_B64	V_MUL_I32_I24	V_CVT_F16_{U16, I16}
V_FMA_{ F16, F32 , F64}	V_ASHRREV_I64	V_MUL_HI_I32_I24	V_CVT_RPI_I32_F32
V_BFI_B32	V_LDEXP_F32	V_MUL_U32_U24	V_CVT_FLR_I32_F32
V_LERP_U8	V_READLANE_B32	V_MUL_HI_U32_U24	V_CVT_OFF_F32_I4
V_ALIGNBIT_B32	V_WRITELANE_B32	V_MIN_{ F16,U16, I16,F32,I32,U32}	V_FRACT_{ F16,F32,F64}
V_ALIGNBYTE_B32	V_BCNT_U32_B32	V_MAX_{ F16,U16, I16,F32,I32,U32}	V_TRUNC_{ F16,F32, F64}
V_MIN3_{F32,I32,U32}	V_MBCNT_LO_U32_B32	V_LSHRREV_{ B16,B32}	V_CEIL_{ F16,F32, F64}
V_MAX3_{F32,I32,U32}	V_MBCNT_HI_U32_B32	V_ASHRREV_{I16,B32}	V_RNDNE_{ F16,F32, F64}
V_MED3_{F32,I32,U32}	V_CVT_PKACCUM_U8_F32	V_LSHLREV_{ B16,B32}	V_FLOOR_{ F16,F32, F64}
V_SAD_{U8, HI_U8, U16, U32}	V_CVT_PKNORM_I16_F32	V_AND_B32	V_EXP_{ F16,F32}
V_CVT_PK_U8_F32	V_CVT_PKNORM_U16_F32	V_OR_B32	V_LOG_{ F16,F32}
V_DIV_FIXUP_{ F16,F32,F64}	V_CVT_PKRTZ_F16_F32	V_XOR_B32	V_RCP_{ F16,F32,F64}
V_DIV_SCALE_{F32,F64}	V_CVT_PK_U16_U32	V_MAC_{ F16,F32}	V_RCP_IFLAG_F32

V_DIV_FMAS_{F32,F64}	V_CVT_PK_I16_I32	V_MADMK_{ F16,F32}	V_RSQ_{ F16,F32, F64}
V_MSAD_U8	V_BFM_B32	V_MADAK_{ F16,F32}	V_SQRT_{ F16,F32,F64}
V_QSAD_PK_U16_U8	V_INTERP_P1_F32	V_CNDMASK_B32	V_SIN_{ F16,F32}
V_MQSAD_PK_U16_U8	V_INTERP_P2_F32	V_LDEXP_F16	V_COS_{ F16,F32}
V_MQSAD_PK_U32_U8	V_INTERP_MOV_F32	MUL_LO_U16	V_NOT_B32
V_TRIG_PREOP_F64	V_INTERP		V_BFREV_B32
V_MAD_{U64_U32, I64_I32}	V_INTERP_P1LL_F16		V_FFBH_U32
	V_INTERP_P1LV_F16		V_FFBH_B32
	V_INTERP_P2_F16		V_FFBH_I32
			V_FREXP_EXP_I32_F64
			V_FREXP_MANT_{ F16,F32,64}
			V_FREXP_EXP_I32_F32
			V_FREXP_EXP_I16_F16
			V_CLREXCP

Table 6.3 lists the compare instruction.

Table 6.3 Compare Operations

VOPC - Compare Ops			
VOPC writes to VCC, VOP3 writes compare result to any SGPR			
V_CMP	I16, I32, I64, U16, U32, U64	F, LT, EQ, LE, GT, LG, GE, T	write VCC
V_CMPX			write VCC and exec
V_CMP	F16, F32, F64	F, LT, EQ, LE, GT, LG, GE, T, O, U, NGE, NLG, NGT, NLE, NEQ, NLT (o = total order, u = unordered, "N" = NaN or normal compare)	write VCC
V_CMPX			write VCC and exec
V_CMP_CLASS	F16, F32, F64	Test for one of: signaling-NaN, quiet-NaN, positive or negative: infinity, normal, sub-normal, zero.	write VCC
V_CMPX_CLASS			write VCC and exec

6.4 Denormalized and Rounding Modes

The shader program has explicit control over the rounding mode applied and the handling of denormalized inputs and results. The MODE register is set using the `S_SETREG` instruction; it has separate bits for controlling the behavior of single- and double-precision floating-point numbers (see Table 6.4).

Table 6.4 MODE Register FP Bits

Field	Bit Position	Description
FP_ROUND	3:0	[1:0] Single-precision round mode. [3:2] Double-precision round mode. Round Modes: 0=nearest even; 1= +infinity; 2= -infinity; 3= toward zero.
FP_DENORM	7:4	[5:4] Single-precision denormal mode. [7:6] Double-precision denormal mode. Denormal modes: 0 = Flush input and output denorms. 1 = Allow input denorms, flush output denorms. 2 = Flush input denorms, allow output denorms. 3 = Allow input and output denorms.

6.5 ALU CLAMP Bit Usage

In GCN Generation 3, the meaning of the “Clamp” bit in the VALU instructions has changed. For `V_CMP` instructions, setting the clamp bit to 1 indicates that the compare signals if a floating point exception occurs. For integer operations, it clamps the result to the largest and smallest representable value. For floating point operations, it clamps the result to the range: [0.0, 1.0].

6.6 VGPR Indexing

VGPR Indexing allows a value stored in the M0 register to act as an index into the VGPRs either for the source or destination registers in VALU instructions.

6.6.1 Indexing Instructions

Table 6.5 describes the instructions which enable, disable and control VGPR indexing.

Table 6.5 VGPR Indexing Instructions

Instruction	Encoding	Sets SCCS?	Operation
S_SET_GPR_IDX_OFF	SOPP	N	Disable VGPR indexing mode. Sets: mode.gpr_idx_en = 0.
S_SET_GPR_IDX_ON	SOPC	N	Enable VGPR indexing, and set the index value and mode from an SGPR. mode.gpr_idx_en = 1 M0[7:0] = S0.u[7:0] M0[15:12] = SIMM4
S_SET_GPR_IDX_IDX	SOP1	N	Set the VGPR index value: M0[7:0] = S0.u[7:0]
S_SET_GPR_IDX_MODE	SOPP	N	Change the VGPR indexing mode, which is stored in M0[15:12]. M0[15:12] = SIMM4

Indexing is enabled and disabled by a bit in the MODE register: gpr_idx_en. When enabled, two fields from M0 are used to determine the index value and what it applies to:

- M0[7:0] holds the unsigned index value, added to selected source or destination VGPR addresses.
- M0[15:12] holds a four-bit mask indicating to which source or destination the index is applied.
 - M0[15] = dest_enable
 - M0[14] = src2_enable
 - M0[13] = src1_enable
 - M0[12] = src0_enable

Indexing only works on VGPR source and destinations, not on inline constants or SGPRs. It is illegal for the index attempt to address VGPRs that are out of range.

6.6.2 Special Cases

This section describes how VGPR indexing is applied to instructions that use source and destination registers in unusual ways. The table below shows which M0 bits control indexing of the sources and destination registers for these special instructions.

Instruction	Microcode Encodes	VALU Receives	Fields to which M0 Index Bits Apply			
			[15] (dst)	[14] (s2)	[13] (s1)	[12] (s0)
v_readlane	sdst = src0, SS1		x	x	x	src0
v_readfirstlane	sdst = func(src0)		x	x	x	src0
v_writelane	dst = func(ss0, ss1)		dst	x	x	x
v_mac_*	dst = src0 * src1 + dst	mad: dst, src0, src1, src2	dst, s2	x	src1	src0
v_madak	dst = src0 * src1 + imm	mad: dst, src0, src1, src2	dst	x	src1	src0
v_madm	dst = S0 * imm + src1	mad: dst, src0, src1, src2	dst	src2	x	src0
v_*sh*_rev	dst = S1 << S0	<shift> (src1, src0)	dst	x	src1	src0
v_cvt_pkaccum	uses dst as src2		dst, s2	x	src1	src0
SDWA (dest preserve, sub-dword mask)	uses dst as src2 for read-mod-write		dst, s2			

where:

src= vector source

SS = scalar source

dst = vector destination

sdst = scalar destination

Chapter 7


Scalar Memory Operations

Scalar Memory Read (SMEM) instructions allow a shader program to load data from memory into SGPRs through the Scalar Data Cache, or write data from SGPRs to memory through the Scalar Data Cache. Instructions can read from 1 to 16 dwords, or write 1 to 4 dwords at a time. Data is read directly into SGPRs with any format conversion.

The scalar unit reads and writes consecutive dwords between memory and the SGPRs. This is intended primarily for loading ALU constants and for indirect T#/S# lookup. No data formatting is supported, nor is byte or short data.

7.1 Microcode Encoding

Scalar memory read instructions are encoded using the SMEM microcode format.



The fields are described in Table 7.1

Table 7.1 SMEM Encoding Field Descriptions

Field	Bits	Description
OP	21:18	Opcode.
IMM	17	Determines how the OFFSET field is interpreted. IMM=1 : Offset is a 20-bit unsigned byte offset to the address. IMM=0 : Offset[6:0] specifies an SGPR or M0 which provides an unsigned byte offset. STORE and ATOMIC instructions cannot use an SGPR: only imm or M0.
GLC	16	Globally Coherent. For loads, controls L1 cache policy: 0=hit_lru, 1=miss_evict. For stores, controls L1 cache bypass: 0=write-combine, 1=write-thru.

Table 7.1 SMEM Encoding Field Descriptions (Cont.)

Field	Bits	Description
SDATA	12:6	SGPRs to return read data to, or to source write-data from. <ul style="list-style-type: none"> • Reads of 2 dwords must have an even SDST-sgpr. • Reads of 4 or more dwords must have their DST-gpr aligned to a multiple of 4. • SDATA must be: SGPR or VCC. Not: exec or m0.
SBASE	5:0	SGPR-pair (SBASE has an implied LSB of zero) which provides a base address, or for BUFFER instructions, a set of 4 SGPRs (4-sgpr aligned) which hold the resource constant. For BUFFER instructions, the only resource fields used are: base, stride, num_records
OFFSET	51:32	An unsigned byte offset, or the address of an SGPR holding the offset. <i>Writes and atomics: M0 or immediate only, not SGPR.</i>

7.2 Operations

7.2.1 S_LOAD_DWORD

These instructions load 1-16 Dwords from memory at the address specified in the SBASE register plus the offset. SBASE holds a 64-bit byte-address.

Memory Address = BASE + OFFSET, truncated to a Dword address.

7.2.2 S_STORE_DWORD

These instructions store 1-4 Dwords of data from SGPRS (starting from SDST) to memory. Addressing is identical to S_LOAD_DWORD. Store instructions cannot use IMM=0, and OFFSET is anything other than M0.

7.2.3 S_BUFFER_LOAD_DWORD, S_BUFFER_STORE_DWORD

These instructions also load 1-16 Dwords from memory or write 1-4 Dwords to memory, but they use a buffer resource (V#, [Chapter 8, "Vector Memory Operations"](#)). The resource provides:

- Base address
- Stride
- Num_records
- *All other buffer resource fields are ignored.*

Memory Address = Base (from V#) + OFFSET, truncated to a Dword address.

The address is clamped if:

- Stride is zero: clamp if (offset >= num_records)
- Stride is non-zero: clamp if (offset > (stride * num_records))

7.2.4 S_DCACHE_INV, S_DCACHE_WB

This instruction invalidates, or does a “write back” of dirty data, for the entire data cache. It does not return anything to SDST.

7.2.5 S_MEM_TIME

This instruction reads a 64-bit clock counter into a pair of SGPRs: SDST and SDST+1.

7.2.6 S_MEM_REALTIME

This instruction reads a 64-bit “real time” counter and returns the value into a pair of SGPRs: SDST and SDST+1. The time value is from a clock for which the frequency is constant (not affected by power modes or core clock frequency changes).

7.3 Dependency Checking

Scalar memory reads and writes can return data out-of-order from how they were issued; they can return partial results at different times when the read crosses two cache lines. The shader program uses the LGKM_CNT counter to determine when the data has been returned to the SDST SGPRs. This is done as follows.

- LGKM_CNT is incremented by 1 for every fetch of a single Dword.
- LGKM_CNT is incremented by 2 for every fetch of two or more Dwords.
- LGKM_CNT is decremented by an equal amount when each instruction completes.

Because the instructions can return out-of-order, the only sensible way to use this counter is to implement `S_WAITCNT 0`; this imposes a wait for all data to return from previous SMEMs before continuing.

7.4 Alignment and Bounds Checking

SDST – The value of SDST must be even for fetches of two Dwords (including S_MEMTIME), or a multiple of four for larger fetches. If this rule is not followed, invalid data can result. If SDST is out-of-range, the instruction is not executed.

SBASE – The value of SBASE must be even for S_BUFFER_LOAD (specifying the address of an SGPR which is a multiple of four). If SBASE is out-of-range, the value from SGPR0 is used.

OFFSET – The value of OFFSET has no alignment restrictions.

Memory Address – If the memory address is out-of-range (clamped), the operation is not performed for any Dwords that are out-of-range.

Chapter 8

Vector Memory Operations

Vector Memory (VMEM) instructions read or write one piece of data separately for each work-item in a wavefront into, or out of, VGPRs. This is in contrast to Scalar Memory instructions, which move a single piece of data that is shared by all threads in the wavefront. All Vector Memory (VM) operations are processed by the texture cache system (level 1 and level 2 caches).

Software initiates a load, store or atomic operation through the texture cache through one of three types of VMEM instructions:

- MTBUF – Memory typed-buffer operations.
- MUBUF – Memory untyped-buffer operations.
- MIMG – Memory image operations.

These instruction types are described by one of three 64-bit microcode formats (see [Section 13.6, “Vector Memory Buffer Instructions,” page 13-50](#) and [Section 13.7, “Vector Memory Image Instruction,” page 13-58](#)). The instruction defines which VGPR(s) supply the addresses for the operation, which VGPRs supply or receive data from the operation, and a series of SGPRs that contain the memory buffer descriptor ($V\#$ or $T\#$). Also, MIMG operations supply a texture sampler from a series of four SGPRs; this sampler defines texel filtering operations to be performed on data read from the image.

8.1 Vector Memory Buffer Instructions

Vector-memory (VM) operations transfer data between the VGPRs and buffer objects in memory through the texture cache (TC). “Vector” means that one or more piece of data is transferred uniquely for every thread in the wavefront, in contrast to scalar memory reads, which transfer only one value that is shared by all threads in the wavefront.

Buffer reads have the option of returning data to VGPRs or directly into LDS.

Examples of buffer objects are vertex buffers, raw buffers, stream-out buffers, and structured buffers.

Buffer objects support both homogenous and heterogeneous data, but no filtering of read-data (no samplers). Buffer instructions are divided into two groups:

- MUBUF – Untyped buffer objects.
 - Data format is specified in the resource constant.
 - Load, store, atomic operations, with or without data format conversion.
- MTBUF – Typed buffer objects.
 - Data format is specified in the instruction.
 - The only operations are Load and Store, both with data format conversion.

Atomic operations take data from VGPRs and combine them arithmetically with data already in memory. Optionally, the value that was in memory before the operation took place can be returned to the shader.

All VM operations use a buffer resource constant (T#) which is a 128-bit value in SGPRs. This constant is sent to the texture cache when the instruction is executed. This constant defines the address and characteristics of the buffer in memory. Typically, these constants are fetched from memory using scalar memory reads prior to executing VM instructions, but these constants also can be generated within the shader.

8.1.1 Simplified Buffer Addressing

The equation in Figure 8.1 shows how the hardware calculates the memory address for a buffer access.

$\text{ADDR} = \underset{\text{T\#}}{\text{Base}} + \underset{\text{SGPR}}{\text{baseOffset}} + \underset{\text{Instr}}{\text{loffset}} + \underset{\text{VGPR}}{\text{Voffset}} + \underset{\text{T\#}}{\text{Stride}} * (\underset{\text{VGPR}}{\text{Vindex}} + \underset{\text{0..63}}{\text{TID}})$ <p style="margin: 0;">Voffset is ignored when instruction bit "OFFEN" == 0 Vindex is ignored when instruction bit "IDXEN" == 0 TID is a constant value (0..63) unique to each thread in the wave. It is ignored when resource bit ADD_TID_ENABLE == 0</p>
--

Figure 8.1 Buffer Address Components

8.1.2 Buffer Instructions

Buffer instructions (MTBUF and MUBUF) allow the shader program to read from, and write to, linear buffers in memory. These operations can operate on data as small as one byte, and up to four Dwords per work-item. Atomic arithmetic operations are provided that can operate on the data values in memory and, optionally, return the value that was in memory before the arithmetic operation was performed.

The D16 instruction variants convert the results to packed 16-bit values. For example, BUFFER_LOAD_FORMAT_D16_XYZW writes two VGPRs. The D16 variants are only available on GCN Generation 3 revision 1 processors.

Table 8.1 Buffer Instructions

Instruction	Description
MTBUF Instructions	
TBUFFER_LOAD_FORMAT_{x,xy,xyz,xyzw} TBUFFER_STORE_FORMAT_{x,xy,xyz,xyzw}	Read from, or write to, a typed buffer object. Also used for a vertex fetch.
MUBUF Instructions	
BUFFER_LOAD_FORMAT_{x,xy,xyz,xyzw} BUFFER_STORE_FORMAT_{x,xy,xyz,xyzw} BUFFER_LOAD_<size> BUFFER_STORE_<size>	Read to, or write from, an untyped buffer object. <size> = byte, ubyte, short, ushort, Dword, Dwordx2, Dwordx3, Dwordx4
BUFFER_ATOMIC_<op> BUFFER_ATOMIC_<op>_ x2	Buffer object atomic operation. Always globally coherent. Operates on 32-bit or 64-bit values (x2 = 64 bits).

Table 8.2 Microcode Formats

Field	Bit Size	Description																												
OP	3 7	MTBUF: Opcode for Typed buffer instructions. MUBUF: Opcode for Untyped buffer instructions.																												
VADDR	8	Address of VGPR to supply first component of address (offset or index). When both index and offset are used, index is in the first VGPR, offset in the second.																												
VDATA	8	Address of VGPR to supply first component of write data or receive first component of read-data.																												
SOFFSET	8	SGPR to supply unsigned byte offset. Must be an SGPR, M0, or inline constant.																												
SRSRC	5	Specifies which SGPR supplies T# (resource constant) in four or eight consecutive SGPRs. This field is missing the two LSBs of the SGPR address, since this address must be aligned to a multiple of four SGPRs.																												
DFMT	4	Data Format of data in memory buffer: <table style="margin-left: 20px; border: none;"> <tr><td>0</td><td>invalid</td><td>8</td><td>10_10_10_2</td></tr> <tr><td>1</td><td>8</td><td>9</td><td>2_10_10_10</td></tr> <tr><td>2</td><td>16</td><td>10</td><td>8_8_8_8</td></tr> <tr><td>3</td><td>8_8</td><td>11</td><td>32_32</td></tr> <tr><td>4</td><td>32</td><td>12</td><td>16_16_16_16</td></tr> <tr><td>5</td><td>16_16</td><td>13</td><td>32_32_32</td></tr> <tr><td>6</td><td>10_11_11</td><td>14</td><td>32_32_32_32</td></tr> </table>	0	invalid	8	10_10_10_2	1	8	9	2_10_10_10	2	16	10	8_8_8_8	3	8_8	11	32_32	4	32	12	16_16_16_16	5	16_16	13	32_32_32	6	10_11_11	14	32_32_32_32
0	invalid	8	10_10_10_2																											
1	8	9	2_10_10_10																											
2	16	10	8_8_8_8																											
3	8_8	11	32_32																											
4	32	12	16_16_16_16																											
5	16_16	13	32_32_32																											
6	10_11_11	14	32_32_32_32																											
NFMT	3	Numeric format of data in memory. <table style="margin-left: 20px; border: none;"> <tr><td>0</td><td>unorm</td></tr> <tr><td>1</td><td>snorm</td></tr> <tr><td>2</td><td>uscaled</td></tr> <tr><td>3</td><td>sscaled</td></tr> <tr><td>4</td><td>uint</td></tr> <tr><td>5</td><td>sint</td></tr> <tr><td>6</td><td>reserved</td></tr> <tr><td>7</td><td>float</td></tr> </table>	0	unorm	1	snorm	2	uscaled	3	sscaled	4	uint	5	sint	6	reserved	7	float												
0	unorm																													
1	snorm																													
2	uscaled																													
3	sscaled																													
4	uint																													
5	sint																													
6	reserved																													
7	float																													
OFFSET	12	Unsigned byte offset.																												
OFFEN	1	1 = Supply an offset from VGPR (VADDR). 0 = Do not (offset = 0).																												

Table 8.2 Microcode Formats (Cont.)

Field	Bit Size	Description
IDXEN	1	1 = Supply an index from VGPR (VADDR). 0 = Do not (index = 0).
GLC	1	Globally Coherent. Controls how reads and writes are handled by the L1 texture cache. READ GLC = 0 Reads can hit on the L1 and persist across wavefronts GLC = 1 Reads always miss the L1 and force fetch to L2. No L1 persistence across waves. WRITE GLC = 0 Writes miss the L1, write through to L2, and persist in L1 across wavefronts. GLC = 1 Writes miss the L1, write through to L2. No persistence across wavefronts. ATOMIC GLC = 0 Previous data value is not returned. No L1 persistence across wavefronts. GLC = 1 Previous data value is returned. No L1 persistence across wavefronts. Note: GLC means “return pre-op value” for atomics.
SLC	1	System Level Coherent. When set, accesses are forced to miss in level 2 texture cache and are coherent with system memory.
TFE	1	Texel Fail Enable for PRT (partially resident textures). When set to 1, fetch can return a NACK that causes a VGPR write into DST+1 (first GPR after all fetch-dest GPRs).
LDS	1	MUBUF-ONLY: 0 = Return read-data to VGPRs. 1 = Return read-data to LDS instead of VGPRs.

8.1.3 VGPR Usage

VGPRs supply address and write-data; also, they can be the destination for return data (the other option is LDS).

Address – Zero, one or two VGPRs are used, depending of the offset-enable (OFFEN) and index-enable (IDXEN) in the instruction word, as shown in Table 8.3.

Table 8.3 Address VGPRs

IDXEN	OFFEN	VGPR _n	VGPR _{n+1}
0	0	<i>nothing</i>	
0	1	uint offset	
1	0	uint index	
1	1	uint index	uint offset

Write Data – *N* consecutive VGPRs, starting at VDATA. The data format specified in the instruction word (NFMT, DFMT for MTBUF, or encoded in the opcode field for MUBUF) determines how many Dwords to write.

Read Data – Same as writes. Data is returned to consecutive GPRs.

Read Data Format – Read data is always 32 bits, based on the data format in the instruction or resource. Float or normalized data is returned as floats; integer formats are returned as integers (signed or unsigned, same type as the memory storage format). Memory reads of data in memory that is 32 or 64 bits do not undergo any format conversion.

Atomics with Return – Data is read out of the VGPR(s) starting at VDATA to supply to the atomic operation. If the atomic returns a value to VGPRs, that data is returned to those same VGPRs starting at VDATA.

8.1.4 Buffer Data

The amount and type of data that is read or written is controlled by the following: data-format (*dfmt*), numeric-format (*nfmt*), destination-component-selects (*dst_sel*), and the opcode. *Dfmt* and *nfmt* can come from the resource, instruction fields, or the opcode itself. *Dst_sel* comes from the resource, but is ignored for many operations.

Table 8.4 Buffer Instructions

Instruction	Data Format	Num Format	DST SEL
TBUFFER_LOAD_FORMAT_*	instruction	instruction	identity
TBUFFER_STORE_FORMAT_*	instruction	instruction	identity
BUFFER_LOAD_<type>	derived	derived	identity
BUFFER_STORE_<type>	derived	derived	identity
BUFFER_LOAD_FORMAT_*	resource	resource	resource
BUFFER_STORE_FORMAT_*	resource	resource	resource
BUFFER_ATOMIC_*	derived	derived	identity

Instruction – The instruction’s *dfmt* and *nfmt* fields are used instead of the resource’s fields.

Data format derived – The data format is derived from the opcode and ignores the resource definition. For example, *buffer_load_ubyte* sets the data-format to 8 and number-format to *uint*.

NOTE: The resource’s data format must not be INVALID; that format has special meaning (unbound resource), and for that case the data format is not replaced by the instruction’s implied data format.

DST_SEL identity – Depending on the number of components in the data-format, this is: X000, XY00, XYZ0, or XYZW.

The MTBUF derives the data format from the instruction. The MUBUF BUFFER_LOAD_FORMAT and BUFFER_STORE_FORMAT instructions use `dst_sel` from the resource; other MUBUF instructions derive data-format from the instruction itself.

8.1.5 Buffer Addressing

A “buffer” is a data structure in memory that is addressed with an “index” and an “offset.” The index points to a particular record of size “stride” bytes, and the offset is the byte-offset within the record. The “stride” comes from the resource, the index from a VGPR (or zero), and the offset from an SGPR or VGPR and also from the instruction itself.

Table 8.5 BUFFER Instruction Fields for Addressing

Field	Size	Description
<code>inst_offset</code>	12	Literal byte offset from the instruction.
<code>inst_idxen</code>	1	Boolean: get index from VGPR when true, or no index when false.
<code>inst_offen</code>	1	Boolean: get offset from VGPR when true, or no offset when false. Note that <code>inst_offset</code> is always present, regardless of this bit.

The “element size” for a buffer instruction is the amount of data the instruction transfers. It is determined by the DFMT field for MTBUF instructions, or from the opcode for MUBUF instructions. It can be 1, 2, 4, 8, or 16 bytes.

Table 8.6 V# Buffer Resource Constant Fields for Addressing

Field	Size	Description
<code>const_base</code>	48	Base address, in bytes, of the buffer resource.
<code>const_stride</code>	14 or 18	Stride of the record in bytes (0 to 16,383 bytes, or 0 to 262,143 bytes). Normally 14 bits, but is extended to 18-bits when: <code>const_add_tid_enable = true</code> used with MUBUF instructions which are not “format” types (or cache invalidate/WB). This is extension intended for use with scratch (private) buffers. If (<code>const_add_tid_enable == true</code> && MUBUF-non-format instruction) <code>const_stride [17:0] = { V#.DFMT[3:0], V#.const_stride[13:0] }</code> else <code>const_stride</code> is 14 bits : { 4'b0, V#.const_stride[13:0] }
<code>const_num_records</code>	32	Number of records in the buffer. In units of: Bytes if: <code>const_stride == 0</code> or <code>const_swizzle_enable == false</code> Otherwise, in units of “stride”.
<code>const_add_tid_enable</code>	1	Boolean. Add <code>thread_ID</code> within the wavefront to the index when true.

Table 8.6 V# Buffer Resource Constant Fields for Addressing

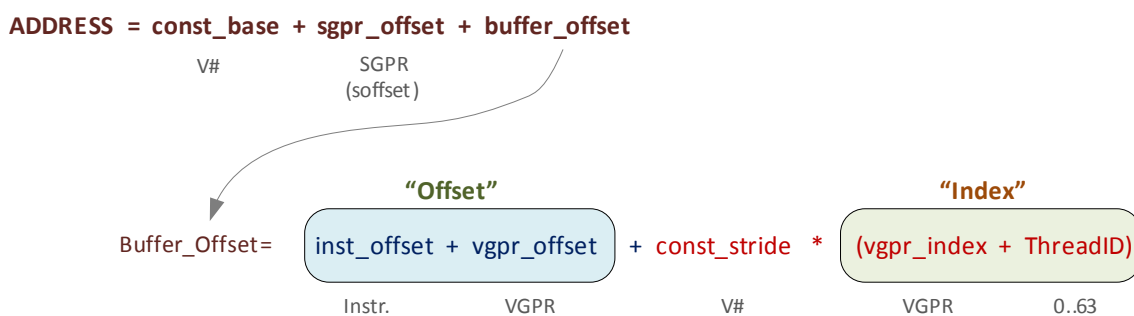
Field	Size	Description
const_swizzle_enable	1	Boolean. Indicates that the surface is swizzled when true.
const_element_size	2	Used only when const_swizzle_en = true. Number of contiguous bytes of a record for a given index (2, 4, 8, or 16 bytes). Must be >= the maximum element size in the structure. const_stride must be an integer multiple of const_element_size.
const_index_stride	2	Used only when const_swizzle_en = true. Number of contiguous indices for a single element (of const_element_size) before switching to the next element. There are 8, 16, 32, or 64 indices.

Table 8.7 Address Components from GPRs

Field	Size	Description
SGPR_offset	32	An unsigned byte-offset to the address. Comes from an SGPR or M0.
VGPR_offset	32	An optional unsigned byte-offset. It is per-thread, and comes from a VGPR.
VGPR_index	32	An optional index value. It is per-thread and comes from a VGPR.

The final buffer memory address is composed of three parts (see Figure 8.2):

- the base address from the buffer resource (V#),
- the offset from the SGPR, and
- a buffer-offset that is calculated differently, depending on whether the buffer is linearly addressed (a simple Array-of-Structures calculation) or is swizzled.



Full equations:

Index = (inst_idxen ? vgpr_index : 0) + (const_add_tid_enable ? thread_id[5:0] : 0)

Offset = (inst_offen ? vgpr_offset : 0) + inst_offset

Figure 8.2 Address Calculation for a Linear Buffer

8.1.5.1 Range Checking

Range checking has changed since Generation 2 with respect to the “const_swizzle_enable=false” case.

It is out of range if:

Case #1: `const_stride == 0 || const_swizzle_enable == false`
const_num_records is in bytes

Out of Range if:

`buffer_offset >= (const_num_records – sgpr_offset)`

Case #2: `const_stride != 0 && const_swizzle_enable == true`
const_num_records is in units of stride

Out of range if either is true:

- `index >= const_num_records`
- `offset >= const_stride`

Notes:

- a. Reads that go out-of-range return zero (except for components with `V#.dst_sel = SEL_1` that return 1).
- b. Writes that are out-of-range do not write anything.
- c. Load/store-format-* instruction and atomics are range-checked “all or nothing” – either entirely in our out.
- d. Load/store-Dword-x{2,3,4} and range-check per component.

8.1.5.2 Swizzled Buffer Addressing

Swizzled addressing rearranges the data in the buffer to provide improved cache locality for arrays of structures. Swizzled addressing also requires Dword-aligned accesses. A single fetch instruction cannot attempt to fetch a unit larger than `const-element-size`. The buffer’s `STRIDE` must be a multiple of `element_size`.

$$\text{Index} = (\text{inst_idxen} ? \text{vgpr_index} : 0) + (\text{const_add_tid_enable} ? \text{thread_id}[5:0] : 0)$$

$$\text{Offset} = (\text{inst_offen} ? \text{vgpr_offset} : 0) + \text{inst_offset}$$


$$\begin{aligned} \text{index_msb} &= \text{index} / \text{const_index_stride} \\ \text{index_lsb} &= \text{index} \% \text{const_index_stride} \\ \text{offset_msb} &= \text{offset} / \text{const_element_size} \\ \text{offset_lsb} &= \text{offset} \% \text{const_element_size} \end{aligned}$$

$$\begin{aligned} \text{buffer_offset} &= (\text{index_msb} * \text{const_stride} + \text{offset_msb} * \\ &\quad \text{const_element_size}) * \text{const_index_stride} + \text{index_lsb} * \\ &\quad \text{const_element_size} + \text{offset_lsb} \end{aligned}$$

$$\text{Final Address} = \text{const_base} + \text{sgpr_offset} + \text{buffer_offset}$$

Remember that the “`sgpr_offset`” is not a part of the “`offset`” term in the above equations. See Figure 8.3.

Original Buffer



Swizzled Buffer

`const_index_stride = 8` // how many consecutive indices to group together
`const_element_size = 4 bytes` // the size of a single element , in bytes


`index_msb = index / const_index_stride`
`index_lsb = index % const_index_stride`
`offset_msb = offset / const_element_size`
`offset_lsb = offset % const_element_size`

$$\text{Buffer_offset} = (\text{index_msb} * \text{const_stride} + \text{offset_msb} * \text{const_element_size}) * \text{const_index_stride} + \text{index_lsb} * \text{const_element_size} + \text{offset_lsb}$$


which simplifies to...

$$\text{Buffer_offset} = (\text{index}/8 * \text{const_stride} + (\text{offset}/4)*4) * 8 + \text{index}\%8 * 4 + \text{offset}\%4$$

Note that because we are dealing with dwords , `offset%4` is always == 0.



If the `const_element_size` had been 8 :



An alternate way to visualize Swizzled Buffers




Figure 8.3 Example of Buffer Swizzling

8.1.5.3 Proposed Uses Cases for Swizzled Addressing

Here are few proposed uses of swizzled addressing in common graphics buffers.

Table 8.8 Address Components from GPRs

	DX11 Raw Uav OpenCL Buffer Object	Dx11 Structured (literal offset)	Dx11 Structured (gpr offset)	Scratch	Ring / stream-out	Const Buffer
inst_vgpr_offset_en	T	F	T	T	T	T
inst_vgpr_index_en	F	T	T	F	F	F
const_stride	na	<api>	<api>	scratchSize	na	na
const_add_tid_enable	F	F	F	T	T	F
const_buffer_swizzle	F	T	T	T	F	F
const_elem_size	na	4	4	4 or 16	na	4
const_index_stride	na	16	16	64		

8.1.6 Alignment

For Dword or larger reads or writes, the two LSBs of the byte-address are ignored, thus forcing Dword alignment.

8.1.7 Buffer Resource

The buffer resource describes the location of a buffer in memory and the format of the data in the buffer. It is specified in four consecutive SGPRs (four aligned SGPRs) and sent to the texture cache with each buffer instruction.

Table 8.9 details the fields that make up the buffer resource descriptor.

Table 8.9 Buffer Resource Descriptor¹

Bits	Size	Name	Description
47:0	48	Base address	Byte address. (In the Northern Islands environment, this was 40 bits.)
61:48	14	Stride	Bytes 0 to 16383
62	1	Cache swizzle	Buffer access. Optionally, swizzle texture cache TC L1 cache banks.
63	1	Swizzle enable	Swizzle AOS according to stride, index_stride, and element_size, else linear (stride * index + offset).
95:64	32	Num_records	In units of stride.
98:96	3	Dst_sel_x	Destination channel select: 0=0, 1=1, 4=R, 5=G, 6=B, 7=A
101:99	3	Dst_sel_y	
104:102	3	Dst_sel_z	
107:105	3	Dst_sel_w	
110:108	3	Num format	Numeric data type (float, int, ...). See instruction encoding for values.
114:111	4	Data format	Number of fields and size of each field. See instruction encoding for values. For MUBUF instructions with "ADD_TID_EN = 1." This field holds Stride [17:14].

Table 8.9 Buffer Resource Descriptor¹ (Cont.)

Bits	Size	Name	Description
116:115	2	Element size	2, 4, 8, or 16 bytes (NI = 4). Used for swizzled buffer addressing.
118:117	2	Index stride	8, 16, 32, or 64 (NI = 16). Used for swizzled buffer addressing.
119	1	Add tid enable	Add thread ID to the index for to calculate the address.
120	1	ATC	0: resource is in GPUVM memory space. 1 = resource is in ATC memory space.
121	1	Hash enable	1 = buffer addresses are hashed for better cache performance.
122	1	Heap	1 = buffer is a heap. out-of-range if offset = 0 or >= num_records.
125:123	3	MTYPE	Memory type - controls cache behavior.
127:126	2	Type	value == 0 for buffer. Overlaps upper two bits of four-bit TYPE field in 128-bit T# resource.

1.A resource set to all zeros acts as an unbound texture or buffer (return 0,0,0,0). Buffer size (in bytes) = (stride==0) ? num_elements : stride * num_elements.

8.1.8 Memory Buffer Load to LDS

The MUBUF instruction format allows reading data from a memory buffer directly into LDS without passing through VGPRs. This is supported for the following subset of MUBUF instructions.

- BUFFER_LOAD_{ubyte, sbyte, ushort, sshort, dword, format_x}.
- It is illegal to set the instruction’s TFE bit for loads to LDS.

LDS_offset = 16-bit unsigned byte offset from M0[15:0].

Mem_offset = 32-bit unsigned byte offset from an SGPR (the SOFFSET SGPR).

idx_vgpr = index value from a VGPR (located at VADDR). (Zero if idxen=0.)

off_vgpr = offset value from a VGPR (located at VADDR or VADDR+1). (Zero if offen=0.)

Figure 8.4 shows the components of the LDS and memory address calculation.

$\text{LDS_ADDR} = \text{LDSbase} + \text{LDS_offset} + (\text{TIDinWave} * 4) + \text{Inst_offset}$							
Alloc	M0[15:0]	0.63	bytes-per-word				
$\text{MEM_ADDR} = \text{Base} + \text{mem_offset} + \text{inst_offset} + \text{off_vgpr} + \text{stride} * (\text{idx_vgpr} + \text{TIDinWave})$							
T#	SGPR	Instr.	VGPR	T#	VGPR	0.63	
(soffset)							

Figure 8.4 Components of Addresses for LDS and Memory

TIDinWave is only added if the resource (T#) has the `ADD_TID_ENABLE` field set to 1. LDS always adds it.

The `MEM_ADDR M#` is in the `VDATA` field; it specifies `M0`.

8.1.8.1 Clamping Rules

Memory address clamping follows the same rules as any other buffer fetch.

LDS address clamping: the return data must not be written outside the LDS space allocated to this wave.

- Set the active-mask to limit buffer reads to those threads that return data to a legal LDS location.
- The `LDSbase (alloc)` is in units of 32 Dwords, as is `LDSsize`.
- `M0[15:0]` is in bytes.

8.1.9 GLC Bit Explained

The GLC bit means different things for loads, stores, and atomic ops.

GLC Meaning for Loads

- For `GLC==0`
 - The load can read data from the GPU L1.
 - Typically, all loads (except load-acquire) use `GLC==0`.
- For `GLC==1`
 - The load intentionally misses the GPU L1 and reads from L2.
If there was a line in the GPU L1 that matched, it is invalidated; L2 is re-read.
 - NOTE: L2 is not re-read for every work-item in the same wave-front for a single load instruction. For example:

```
b=uav[N+tid] // assume this is a byte read w/ glc==1 and N is aligned to 64B
```

In the above op, the first Tid of the wavefront brings in the line from L2 or beyond, and all 63 of the other Tids read from same 64 B cache line in the L1.

GLC Meaning for Stores

- For `GLC==0`

This causes a write-combine across work-items of the wavefront store op; dirtied lines are written to the L2 automatically.

 - If the store operation dirtied all bytes of the 64 B line, it is left clean and valid in the L1; subsequent accesses to the cache are allowed to hit on this cache line.

- Else do not leave write-combined lines in L1.
- For GLC==1
 - Same as GLC==0, except the write-combined lines are not left in the line, even if all bytes are dirtied.

Atomic

- For GLC == 0
 - No return data (this is “write-only” atomic op).
- For GLC == 1
 - Returns previous value in memory (before the atomic operation).

8.2 Vector Memory (VM) Image Instructions

Vector Memory (VM) operations transfer data between the VGPRs and memory through the texture cache (TC). Vector means the transfer of one or more pieces of data uniquely for every work-item in the wavefront. This is in contrast to scalar memory reads, which transfer only one value that is shared by all work-items in the wavefront.

Examples of image objects are texture maps and typed surfaces.

Image objects are accessed using from one to four dimensional addresses; they are composed of homogenous data of one to four elements. These image objects are read from, or written to, using `IMAGE_*` or `SAMPLE_*` instructions, all of which use the MIMG instruction format. `IMAGE_LOAD` instructions read an element from the image buffer directly into VGPRs, and `SAMPLE` instructions use sampler constants (S#) and apply filtering to the data after it is read. `IMAGE_ATOMIC` instructions combine data from VGPRs with data already in memory, and optionally return the value that was in memory before the operation.

All VM operations use an image resource constant (T#) that is a 128- or 256-bit value in SGPRs. This constant is sent to the texture cache when the instruction is executed. This constant defines the address, data format, and characteristics of the surface in memory. Some image instructions also use a sampler constant that is a 128-bit constant in SGPRs. Typically, these constants are fetched from memory using scalar memory reads prior to executing VM instructions, but these constants can also be generated within the shader.

Texture fetch instructions have a data mask (DMASK) field. DMASK specifies how many data components it receives. If DMASK is less than the number of components in the texture, the texture unit only sends DMASK components, starting with R, then G, B, and A. If DMASK specifies more than the texture format specifies, the shader receives zero for the missing components.

8.2.1 Image Instructions

This section describes the image instruction set, and the microcode fields available to those instructions.

Table 8.10 Image Instructions

MIMG Instruction	Description
SAMPLE_*	Read and filter data from a image object.
IMAGE_LOAD_<op>	Read data from an image object using one of the following: image_load, image_load_mip, image_load_{pck, pck_sgn, mip_pck, mip_pck_sgn}.
IMAGE_STORE IMAGE_STORE_MIP	Store data to an image object. Store data to a specific mipmap level.
IMAGE_ATOMIC_<op>	Image atomic operation, which is one of the following: swap, cmpswap, add, sub, rsub, {u,s}{min,max}, and, or, xor, inc, dec, fcmpswap, fmin, fmax.

Table 8.11 Instruction Fields

Instruction	Bit Size	Description
OP	7	Opcode.
VADDR	8	Address of VGPR to supply first component of address.
VDATA	8	Address of VGPR to supply first component of write data or receive first component of read-data.
SSAMP	5	SGPR to supply S# (sampler constant) in four consecutive SGPRs. Missing two LSBs of SGPR-address since must be aligned to a multiple of four SGPRs.
SRSRC	5	SGPR to supply T# (resource constant) in four or eight consecutive SGPRs. Missing two LSBs of SGPR-address since must be aligned to a multiple of four SGPRs.
UNRM	1	Force address to be un-normalized regardless of T#. Must be set to 1 for image stores and atomics.
R128	1	Texture resource size: 1 = 128 bits, 0 = 256 bits.
DA	1	Shader declared an array resource to be used with this fetch. When 1, the shader provides an array-index with the instruction. When 0, no array index is provided.
DMASK	4	Data VGPR enable mask: one to four consecutive VGPRs. Reads: defines which components are returned. 0 = red, 1 = green, 2 = blue, 3 = alpha Writes: defines which components are written with data from VGPRs (missing components get 0). Enabled components come from consecutive VGPRs. For example: DMASK=1001 : Red is in VGPRn and alpha in VGPRn+1. If DMASK==0, the TA overrides the data format to "invalid," and forces dst_sels to return 0.

Table 8.11 Instruction Fields (Cont.)

Instruction	Bit Size	Description
GLC	1	Globally Coherent. Controls how reads and writes are handled by the L1 texture cache. READ GLC = 0 Reads can hit on the L1 and persist across waves. GLC = 1 Reads always miss the L1 and force fetch to L2. No L1 persistence across waves. WRITE GLC = 0 Writes miss the L1, write through to L2, and persist in L1 across wavefronts. GLC = 1 Writes miss the L1, write through to L2. No persistence across wavefronts. ATOMIC GLC = 0 Previous data value is not returned. No L1 persistence across wavefronts. GLC = 1 Previous data value is returned. No L1 persistence across wavefronts.
SLC	1	System Level Coherent. When set, accesses are forced to miss in level 2 texture cache and are coherent with system memory.
TFE	1	Texel Fail Enable for PRT (partially resident textures). When set, a fetch can return a NACK, which causes a VGPR write into DST+1 (first GPR after all fetch-dest GPRs).
LWE	1	Force data to be un-normalized, regardless of T#.

8.2.2 Image Opcodes with No Sampler

For image opcodes with no sampler, all VGPR address values are taken as uint. For cubemaps, face_id = slice * 6 + face.

Table 8.12 shows the contents of address VGPRs for the various image opcodes.

Table 8.12 Image Opcodes with No Sampler

Image Opcode (Resource w/o Sampler)	Acnt	dim	VGPRn	VGPRn+1	VGPRn+2	VGPRn+3
get_resinfo	0	Any	mipid			
load / store / atomics	0	1D	x			
	1	1D Array	x	slice		
	1	2D	x	y		
	2	2D MSAA	x	y	fragid	
	2	2D Array	x	y	slice	
	3	2D Array MSAA	x	y	slice	fragid
	2	3D	x	y	z	
2	Cube	x	y	face_id		

Table 8.12 Image Opcodes with No Sampler (Cont.)

Image Opcode (Resource w/o Sampler)	Acnt	dim	VGPRn	VGPRn+1	VGPRn+2	VGPRn+3
load_mip / store_mip	1	1D	x	mipid		
	2	1D Array	x	slice	mipid	
	2	2D	x	y	mipid	
	3	2D Array	x	y	slice	mipid
	3	3D	x	y	z	mipid
	3	Cube	x	y	face_id	mipid

8.2.3 Image Opcodes with Sampler

For image opcodes with a sampler, all VGPR address values are taken as float. For cubemaps, face_id = slice * 8 + face.

Certain sample and gather opcodes require additional values from VGPRs beyond what is shown in Table 8.13. These values are: offset, bias, z-compare, and gradients. See Section 8.2.4, “VGPR Usage,” page 8-19, for details.

Table 8.13 Image Opcodes with Sampler

Image Opcode (w/ Sampler)	Acnt	dim	VGPRn	VGPRn+1	VGPRn+2	VGPRn+3
sample ¹	0	1D	x			
	1	1D Array	x	slice		
	1	2D	x	y		
	2	2D interlaced	x	y	field	
	2	2D Array	x	y	slice	
	2	3D	x	y	z	
	2	Cube	x	y	face_id	
sample _i ²	1	1D	x	lod		
	2	1D Array	x	slice	lod	
	2	2D	x	y	lod	
	3	2D interlaced	x	y	field	lod
	3	2D Array	x	y	slice	lod
	3	3D	x	y	z	lod
	3	Cube	x	y	face_id	lod

Table 8.13 Image Opcodes with Sampler (Cont.)

Image Opcode (w/ Sampler)	Acnt	dim	VGPRn	VGPRn+1	VGPRn+2	VGPRn+3
sample_cl ³	1	1D	x	clamp		
	2	1D Array	x	slice	clamp	
	2	2D	x	y	clamp	
	3	2D interlaced	x	y	field	clamp
	3	2D Array	x	y	slice	clamp
	3	3D	x	y	z	clamp
	3	Cube	x	y	face_id	clamp
gather4 ⁴	1	2D	x	y		
	2	2D interlaced	x	y	field	
	2	2D Array	x	y	slice	
	2	Cube	x	y	face_id	
gather4_l	2	2D	x	y	lod	
	3	2D interlaced	x	y	field	lod
	3	2D Array	x	y	slice	lod
	3	Cube	x	y	face_id	lod
gather4_cl	2	2D	x	y	clamp	
	3	2D interlaced	x	y	field	clamp
	3	2D Array	x	y	slice	clamp
	3	Cube	x	y	face_id	clamp

1. sample includes sample, sample_d, sample_b, sample_lz, sample_c, sample_c_d, sample_c_b, sample_c_lz, and getlod
2. sample_l includes sample_l and sample_c_l.
3. sample_cl includes sample_cl, sample_d_cl, sample_b_cl, sample_c_cl, sample_c_d_cl, and sample_c_b_cl.
4. gather4 includes gather4, gather4_lz, gather4_c, and gather4_c_lz.

Table 8.14 lists and briefly describes the legal suffixes for image instructions.

Table 8.14 Sample Instruction Suffix Key

Suffix	Meaning	Extra Addresses	Description
_L	LOD	-	LOD is used instead of TA computed LOD.
_B	LOD BIAS	1: lod bias	Add this BIAS to the LOD TA computes.
_CL	LOD CLAMP	-	Clamp the LOD to be no larger than this value.
_D	Derivative	2,4 or 6: slopes	Send dx/dv, dx/dy, etc. slopes to TA for it to used in LOD computation.
_CD	Coarse Derivative		Send dx/dv, dx/dy, etc. slopes to TA for it to used in LOD computation.

Table 8.14 Sample Instruction Suffix Key (Cont.)

Suffix	Meaning	Extra Addresses	Description
_LZ	Level 0	-	Force use of MIP level 0.
_C	PCF	1: z-comp	Percentage closer filtering.
_O	Offset	1: offsets	Send X, Y, Z integer offsets (packed into 1 Dword) to offset XYZ address.

8.2.4 VGPR Usage

- Address: The address consists of up to four parts:
 { offset } { bias } { z-compare } { derivative } { body }

These are all packed into consecutive VGPRs.

- Offset: `SAMPLE*_O_*`, `GATHER*_O_*`
 One Dword of offset_xyz. The offsets are six-bit signed integers: X=[5:0], Y=[13:8], and Z=[21:16].
- Bias: `SAMPLE*_B_*`, `GATHER*_B_*`. One Dword float.
- Z-compare: `SAMPLE*_C_*`, `GATHER*_C_*`. One Dword.
- Derivatives (sample_d, sample_cd): 2, 4, or 6 Dwords, packed one Dword per derivative as:

Image Dim	VGPR N	N+1	N+2	N+3	N+4	N+5
1D	DX/DH	DX/DV	-	-	-	-
2D	dx/dh	DY/DH	DX/DV	DY/DV	-	--
3D	dx/dh	DY/DH	DZ/DH	DX/DV	DY/DV	DZ/DV

- Body: One to four Dwords, as defined by [Table 8.13](#).
 Address components are X,Y,Z,W with X in VGPR_M, Y in VGPR_M+1, etc.
- Data: Written from, or returned to, one to four consecutive VGPRs. The amount of data read or written is determined by the DMASK field of the instruction.
- Reads: DMASK specifies which elements of the resource are returned to consecutive VGPRs. The texture system reads data from memory and based on the data format expands it to a canonical RGBA form, filling in zero or one for missing components. Then, DMASK is applied, and only those components selected are returned to the shader.
- Writes: When writing an image object, it is only possible to write an entire element (all components), not just individual components. The components come from consecutive VGPRs, and the texture system fills in the value zero for any missing components of the image's data format; it ignores any values that are not part of the stored data format. For example, if the DMASK=1001, the shader sends Red from VGPR_N, and Alpha from VGPR_N+1, to the texture unit. If the image object is RGB, the texel is overwritten with Red from

the VGPR_N, Green and Blue set to zero, and Alpha from the shader ignored.

- **Atomics:** Image atomic operations are supported only on 32- and 64-bit-per-pixel surfaces. The surface data format is specified in the resource constant. Atomic operations treat the element as a single component of 32- or 64-bits. For atomic operations, DMASK is set to the number of VGPRs (Dwords) to send to the texture unit.

DMASK legal values for atomic image operations: no other values of DMASK are legal.

0x1 = 32-bit atomics except cmpswap.

0x3 = 32-bit atomic cmpswap.

0x3 = 64-bit atomics except cmpswap.

0xf = 64-bit atomic cmpswap.

- **Atomics with Return:** Data is read out of the VGPR(s), starting at VDATA, to supply to the atomic operation. If the atomic returns a value to VGPRs, that data is returned to those same VGPRs starting at VDATA.

8.2.5 Image Resource

The image resource (also referred to as T#) defines the location of the image buffer in memory, its dimensions, tiling, and data format. These resources are stored in four or eight consecutive SGPRs and are read by MIMG instructions.

Table 8.15 Image Resource Definition

Bits	Size	Name	Comments
128-bit Resource: 1D-tex, 2d-tex, 2d-msaa (multi-sample auto-aliasing)			
39:0	40	base address	256-byte aligned. Also used for fmask-ptr.
51:40	12	min lod	4.8 (four uint bits, eight fraction bits) format.
57:52	6	data format	Number of comps, number of bits/comp.
61:58	4	num format	Numeric format.
63:62	2	MTYPE[1:0]	Memory type - controls cache behavior.
77:64	14	width	
91:78	14	height	
94:92	3	perf modulation	Scales sampler's perf_z, perf_mip, aniso_bias, lod_bias_sec.
95	1	interlaced	
98:96	3	dst_sel_x	0 = 0, 1 = 1, 4 = R, 5 = G, 6 = B, 7 = A.
101:99	3	dst_sel_y	
104:102	3	dst_sel_z	
107:105	3	dst_sel_w	
111:108	4	base level	
115:112	4	last level	For msaa, holds number of samples

Table 8.15 Image Resource Definition (Cont.)

Bits	Size	Name	Comments
120:116	5	Tiling index	Lookuptable: 32 x 16 bank_width[2], bank_height[2], num_banks[2], tile_split[2], macro_tile_aspect[2], micro_tile_mode[2], array_mode[4].
121	1	pow2pad	Memory footprint is padded to pow2 dimensions
122	1	MTYPE[2]	Bit 2 of the MTYPE field.
123	1	ATC	0 = image is in GPUVM memory; 1 = image is in ATC memory.
127:124	4	type	0 = buf, 8 = 1d, 9 = 2d, 10 = 3d, 11 = cube, 12 = 1d-array, 13 = 2d-array, 14 = 2d-msaa, 15 = 2d-msaa-array. 1-7 are reserved.
256-bit Resource: 1d-array, 2d-array, 3d, cubemap, MSAA			
140:128	13	depth	
154:141	14	pitch	In texel units.
159:155	5	<i>unused</i>	
172:160	13	base array	
185:173	13	last array	
191:186	6	<i>unused</i>	
203:192	12	min_lod_warn	feedback trigger for lod
255:204	52	<i>unused</i>	

All image resource view descriptors (T#'s) are written by the driver as 256 bits. It is permissible to use only the first 128 bits when a simple 1D or 2D (not an array) is bound. This is specified in the MIMG R128 instruction field.

The MIMG-format instructions have a DeclareArray (DA) bit that reflects whether the shader was expecting an array-texture or simple texture to be bound. When DA is zero, the hardware does not send an array index to the texture cache. If the texture map was indexed, the hardware supplies an index value of zero. Indices sent for non-indexed texture maps are ignored.

8.2.6 Sampler Resource

The sampler resource (also referred to as S#) defines what operations to perform on texture map data read by “sample” instructions. These are primarily address clamping and filter options. Sampler resources are defined in four consecutive SGPRs and are supplied to the texture cache with every sample instruction.

Table 8.16 Sampler Resource Definition

Bits	Size	Name	Description
2:0	3	clamp x	Clamp/wrap mode.
5:3	3	clamp y	
8:6	3	clamp z	

Table 8.16 Sampler Resource Definition (Cont.)

Bits	Size	Name	Description
11:9	3	max aniso ratio	
14:12	3	depth compare func	
15	1	force unnormalized	Force address cords to be unorm.
18:16	3	aniso threshold	
19	1	mc coord trunc	
20	1	force degamma	
26:21	6	aniso bias	u1.5.
27	1	trunc coord	
28	1	disable cube wrap	
30:29	2	filter_mode	Normal lerp, min, or max filter.
31	1	compat_mode	1 = new mode; 0 = legacy
43:32	12	min lod	u4.8.
55:44	12	max lod	u4.8.
59:56	4	perf_mip	
63:60	4	perf z	
77:64	14	lod bias	s5.8.
83:78	6	lod bias sec	s1.4.
85:84	2	xy mag filter	Magnification filter.
87:86	2	xy min filter	Minification filter.
89:88	2	z filter	
91:90	2	mip filter	
92	1	mip_point_preclamp	When mipfilter = point, add 0.5 before clamping.
93	1	disable_lsb_ceil	Disable ceiling logic in filter (rounds up).
94	1	Filter_Prec_Fix	
95	1	<i>Aniso_overnde</i>	Disable Aniso filtering if base_level = last_level
107:96	12	border color ptr	
125:108	18	<i>unused</i>	
127:126	2	border color type	Opaque-black, transparent-black, white, use border color ptr.

8.2.7 Data Formats

Data formats 0-15 are available to buffer resources, and all formats are available to image formats. [Table 8.17](#) details all the data formats that can be used by image and buffer resources.

Table 8.17 Data and Image Formats

data_format							shader_num_format						
value	encode	buffer_r	buffer_w	image_r	image_w	MRT (CB)	value	encode	buffer_r	buffer_w	image_r	image_w	
0	invalid	yes	yes	yes	yes	yes	0	unorm	yes	yes	yes	yes	
1	8	yes	yes	yes	yes	yes	1	snorm	yes	yes	yes	yes	
2	16	yes	yes	yes	yes	yes	2	uscaled	yes	no	yes	no	
3	8_8	yes	yes	yes	yes	yes	3	sscaled	yes	no	yes	no	
4	32	yes	yes	yes	yes	yes	4	uint	yes	yes	yes	yes	
5	16_16	yes	yes	yes	yes	yes	5	sint	yes	yes	yes	yes	
6	10_11_11	yes	yes	yes	yes	yes	6	reserved					
7	11_11_10	yes	yes	yes	yes	yes	7	float	yes	yes	yes	yes	
8	10_10_10_2	yes	yes	yes	yes	yes	8	reserved					
9	2_10_10_10	yes	yes	yes	yes	yes	9	srgb	no	no	yes	no	
10	8_8_8_8	yes	yes	yes	yes	yes	10-15	reserved					
11	32_32	yes	yes	yes	yes	yes							
12	16_16_16_16	yes	yes	yes	yes	yes							
13	32_32_32	yes	yes	yes	no	no							
14	32_32_32_32	yes	yes	yes	yes	yes							
15	reserved												
16	5_6_5	no	no	yes	yes	yes							
17	1_5_5_5	no	no	yes	yes	yes							
18	5_5_5_1	no	no	yes	yes	yes							
19	4_4_4_4	no	no	yes	yes	yes							
20	8_24	no	no	yes	no	yes							
21	24_8	no	no	yes	no	yes							
22	X24_8_32	no	no	yes	no	yes							
23-31	reserved												
32	GB_GR	no	no	yes	no	no							
33	BG_RG	no	no	yes	no	no							
34	5_9_9_9	no	no	yes	no	no							
35	BC1	no	no	yes	no	no							
36	BC2	no	no	yes	no	no							
37	BC3	no	no	yes	no	no							
38	BC4	no	no	yes	no	no							
39	BC5	no	no	yes	no	no							

Table 8.17 Data and Image Formats (Cont.)

data_format							shader_num_format					
value	encode	buffer r	buffer w	image r	image w	MRT (CB)	value	encode	buffer r	buffer w	image r	image w
40	BC6	no	no	yes	no	no						
41	BC7	no	no	yes	no	no						
42-46	reserved											
47	FMASK_8_1	no	no	yes	yes	no						
48	FMASK_8_2	no	no	yes	yes	no						
49	FMASK_8_4	no	no	yes	yes	no						
50	FMASK_16_1	no	no	yes	yes	no						
51	FMASK_16_2	no	no	yes	yes	no						
52	FMASK_32_2	no	no	yes	yes	no						
53	FMASK_32_4	no	no	yes	yes	no						
54	FMASK_32_8	no	no	yes	yes	no						
55	FMASK_64_4	no	no	yes	yes	no						
56	FMASK_64_8	no	no	yes	yes	no						
57	4_4	no	no	yes	no	no						
58	6_5_5	no	no	yes	no	no						
59	1	no	no	yes	no	no						
60	1_REVERSE D	no	no	yes	no	no						
61	32_AS_8	no	no	yes	no	no						
62	32_AS_8_8	no	no	yes	no	no						
63	32_AS_32_3 2_32_32	no	no	yes	no	no						

8.2.8 Vector Memory Instruction Data Dependencies

When a VM instruction is issued, the address is immediately read out of VGPRs and sent to the texture cache. Any texture or buffer resources and samplers are also sent immediately. However, write-data is not immediately sent to the texture cache.

The shader developer's responsibility to avoid data hazards associated with VMEM instructions include waiting for VMEM read instruction completion before reading data fetched from the TC (VMCNT).

This is explained in [Section 4.4, "Data Dependency Resolution," page 4-2](#).

Chapter 9

Flat Memory Instructions

Flat Memory instructions read, or write, one piece of data into, or out of, VGPRs; they do this separately for each work-item in a wavefront. Unlike buffer or image instructions, Flat instructions do not use a resource constant to define the base address of a surface. Instead, Flat instructions use a single flat address from the VGPR; this addresses memory as a single flat memory space. This memory space includes video memory, system memory, LDS memory, and scratch (private) memory. It does not include GDS memory. Parts of the flat memory space can not map to any real memory, and accessing these regions generates a memory-violation error. The determination of the memory space to which an address maps is controlled by a set of “memory aperture” base and size registers.

9.1 Flat Memory Instructions

Flat memory instructions let the kernel read or write data in memory, or perform atomic operations on data already in memory. These operations occur through the texture L2 cache. The instruction declares which VGPR holds the address (either 32- or 64-bit, depending on the memory configuration), the VGPR which

sends and the VGPR which receives data. Flat instructions also use M0 as described in Table 9.1.

Table 9.1 Flat Microcode Formats

Field	Bit Size	Description
OP	7	FLAT_LOAD_UBYTE FLAT_STORE_BYTE FLAT_ATOMIC_SWAP FLAT_ATOMIC_SWAP_X2
		FLAT_LOAD_SBYTE FLAT_ATOMIC_CMPSWAP FLAT_ATOMIC_CMPSWAP_X2
		FLAT_LOAD_USHORT FLAT_STORE_SHORT FLAT_ATOMIC_ADD FLAT_ATOMIC_ADD_X2
		FLAT_LOAD_SSHORT FLAT_ATOMIC_SUB FLAT_ATOMIC_SUB_X2
		FLAT_LOAD_DWORD FLAT_STORE_DWORD FLAT_ATOMIC_SMIN FLAT_ATOMIC_SMIN_X2
		FLAT_LOAD_DWORDX2 FLAT_STORE_DWORDX2 FLAT_ATOMIC_UMIN FLAT_ATOMIC_UMIN_X2
		FLAT_LOAD_DWORDX3 FLAT_STORE_DWORDX3 FLAT_ATOMIC_SMAX FLAT_ATOMIC_SMAX_X2
		FLAT_LOAD_DWORDX4 FLAT_STORE_DWORDX4 FLAT_ATOMIC_UMAX FLAT_ATOMIC_UMAX_X2
		FLAT_ATOMIC_AND FLAT_ATOMIC_AND_X2
		FLAT_ATOMIC_OR FLAT_ATOMIC_OR_X2
		FLAT_ATOMIC_XOR FLAT_ATOMIC_XOR_X2
		FLAT_ATOMIC_INC FLAT_ATOMIC_INC_X2
		FLAT_ATOMIC_DEC FLAT_ATOMIC_DEC_X2
ADDR	8	VGPR which holds the address. For 64-bit addresses, ADDR has the LSBs, and ADDR+1 has the MSBs.
DATA	8	VGPR which holds the first Dword of data. Instructions can use 0-4 Dwords.
VDST	8	VGPR destination for data returned to the kernel, either from LOADs or Atomics with GLC=1 (return pre-op value).
SLC	1	System Level Coherent. Used in conjunction with GLC and MTYPE to determine cache policies.
GLC	1	Global Level Coherent. For Atomics, GLC: 1 means return pre-op value, 0 means do not return pre-op value.
TFE	1	Texel Fail Enable for PRT (Partially Resident Textures). When set, fetch can return a NACK, which causes a VGPR write into DST+1 (first GPR after all fetch-dest gprs).
(M0)	32	Implied use of M0. M0[16:0] contains the byte-size of the LDS segment. This is used to clamp the final address.

9.2 Instructions

The FLAT instruction set is nearly identical to the Buffer instruction set, but without the FORMAT reads and writes. Unlike Buffer instructions, FLAT instructions cannot return data directly to LDS, but only to VGPRs.

FLAT instructions do not use a resource constant (V#) or sampler (S#); however, they do require a special SGPR-pair to hold scratch-space information in case any threads' address resolves to scratch space. See [Section 9.6, "Scratch Space \(Private\)," page 9-4](#).

Internally, FLAT instruction are executed as both an LDS and a Buffer instruction; so, they increment both VM_CNT and LGKM_CNT and are not considered done until both have been decremented. There is no way beforehand to determine whether a FLAT instruction uses only LDS or TA memory space.

9.2.1 Ordering

Flat instructions can complete out of order with each other. If one flat instruction finds all of its data in Texture cache, and the next finds all of its data in LDS, the second instruction might complete first. If the two fetches return data to the same VGPR, the result are unknown.

9.2.2 Important Timing Consideration

Since the data for a FLAT load can come from either LDS or the texture cache, and because these units have different latencies, there is a potential race condition with respect to the `VM_CNT` and `LGKM_CNT` counters. Because of this, the only sensible `S_WAITCNT` value to use after FLAT instructions is zero.

9.3 Addressing

FLAT instructions support both 64- and 32-bit addressing. The address size is set using a mode register (`PTR32`), and a local copy of the value is stored per wave.

The addresses for the aperture check differ in 32- and 64-bit mode; however, this is not covered here.

64-bit addresses are stored with the LSBs in the VGPR at `ADDR`, and the MSBs in the VGPR at `ADDR+1`.

For scratch space, the TA takes the address from the VGPR and does the following.

```
Address = VGPR[addr] + TID_in_wave * Size
          + SH_HIDDEN_PRIVATE_BASE_VMID
          - "private aperture base" (in SH_MEM_BASES)
          + offset (from flat_scratch)
```

9.4 Memory Error Checking

Both TA and LDS can report that an error occurred due to a bad address. This can occur for the following reasons:

- invalid address (outside any aperture)
- write to read-only surface
- misaligned data
- out-of-range address:
 - LDS access with an address outside the range: `[0, MIN(M0, LDS_SIZE)-1]`
 - Scratch access with an address outside the range: `[0, scratch-size -1]`
 - Heap address outside of legal range

The policy for threads with bad addresses is: writes outside this range do not write a value, and reads return zero.

Addressing errors from either LDS or TA are returned on their respective “instruction done” busses as `MEM_VIOL`. This sets the wave’s `MEM_VIOL` `TrapStatus` bit and causes an exception (trap) if the corresponding `EXCPEN` bit is set.

9.5 Data

FLAT instructions can use zero to four consecutive Dwords of data in VGPRs and/or memory. The `DATA` field determines which VGPR(s) supply source data (if any), and the `VDST` VGPRs hold return data (if any). No data-format conversion is done.

9.6 Scratch Space (Private)

Scratch (thread-private memory) is an area of memory defined by the aperture registers. When an address falls in scratch space, additional address computation is automatically performed by the hardware. The kernel must provide additional information for this computation to occur in the form of the `FLAT_SCRATCH` register.

The wavefront must supply the scratch size and offset (for space allocated to this wave) with every FLAT request. This is stored in a fixed SGPR location (`FLAT_SCRATCH`): `N_SGPRS-5` and `N_SGPRS-6`, as:

```
{ 8'h0, Offset[31:8], 13'h0, Size[18:0] }
```

- ◇ Offset is in units of 256-bytes (hence the missing eight LSBs)
- ◇ Size is the per-thread scratch size, in bytes.

These SGPRs are automatically sent with every FLAT request.

It is the responsibility of the kernel to initialize this SGPR-pair.

Note that in FSA32, only `SIZE[15:0]` are considered (`[18:16]` are ignored).

Chapter 10

Data Share Operations

Local data share (LDS) is a very low-latency, RAM scratchpad for temporary data with at least one order of magnitude higher effective bandwidth than direct, uncached global memory. It permits sharing of data between work-items in a work-group, as well as holding parameters for pixel shader parameter interpolation. Unlike read-only caches, the LDS permits high-speed write-to-read re-use of the memory space (full gather/read/load and scatter/write/store operations).

10.1 Overview

Figure 10.1 shows the conceptual framework of the LDS is integration into the memory of AMD GPUs using OpenCL.

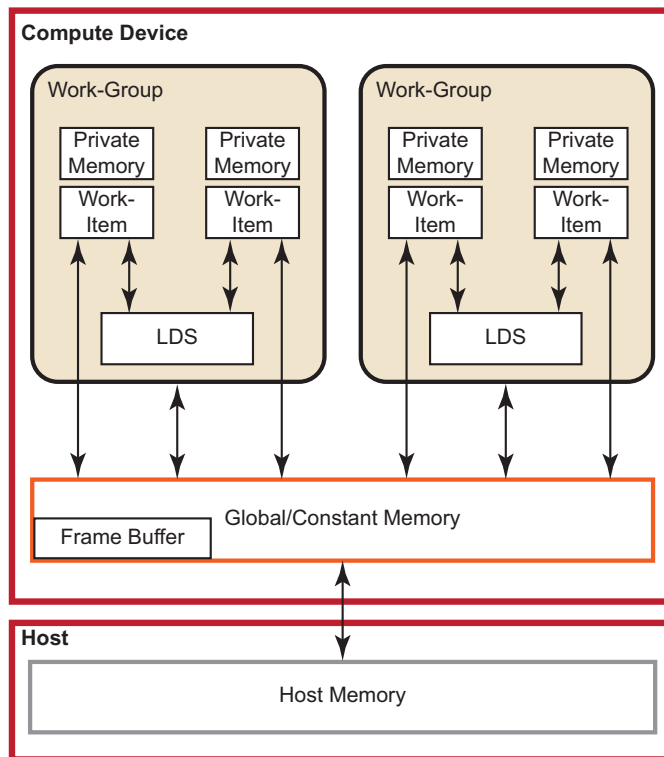


Figure 10.1 High-Level Memory Configuration

Physically located on-chip, directly next to the ALUs, the LDS is approximately one order of magnitude faster than global memory (assuming no bank conflicts).

There are 32 kB memory per compute unit, segmented into 32 or 16 banks (depending on the GPU type) of 1 k dwords (for 32 banks) or 2 k dwords (for 16 banks). Each bank is a 256x32 two-port RAM (1R/1W per clock cycle). Dwords are placed in the banks serially, but all banks can execute a store or load simultaneously. One work-group can request up to 32 kB memory. Reads across wavefront are dispatched over four cycles in waterfall.

The high bandwidth of the LDS memory is achieved not only through its proximity to the ALUs, but also through simultaneous access to its memory banks. Thus, it is possible to concurrently execute 32 write or read instructions, each nominally 32-bits; extended instructions, read2/write2, can be 64-bits each. If, however, more than one access attempt is made to the same bank at the same time, a bank conflict occurs. In this case, for indexed and atomic operations, hardware prevents the attempted concurrent accesses to the same bank by turning them into serial accesses. This decreases the effective bandwidth of the LDS. For maximum throughput (optimal efficiency), therefore, it is important to avoid bank conflicts. A knowledge of request scheduling and address mapping is key to achieving this.

10.2 Dataflow in Memory Hierarchy

Figure 10.2 is a conceptual diagram of the dataflow within the memory structure.

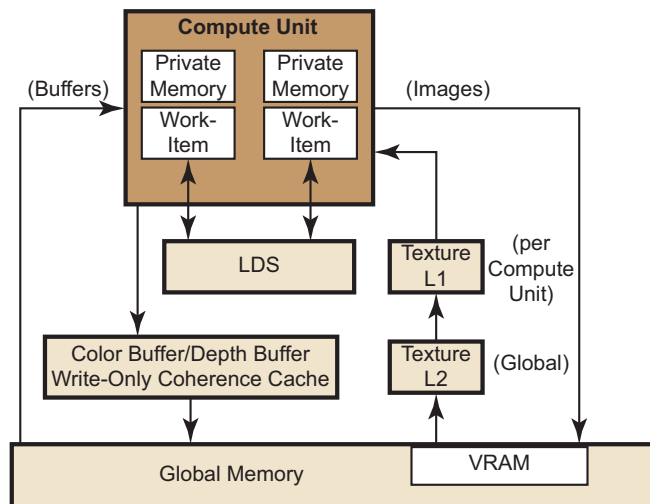


Figure 10.2 Memory Hierarchy Dataflow

To load data into LDS from global memory, it is read from global memory and placed into the work-item's registers; then, a store is performed to LDS. Similarly, to store data into global memory, data is read from LDS and placed into the work-item's registers, then placed into global memory. To make effective use of the LDS, an algorithm must perform many operations on what is transferred between global memory and LDS. It also is possible to load data from a memory buffer directly into LDS, bypassing VGPRs.

LDS atomics are performed in the LDS hardware. (Thus, although ALUs are not directly used for these operations, latency is incurred by the LDS executing this function.) If the algorithm does not require write-to-read reuse (the data is read only), it usually is better to use the image dataflow (see right side of Figure 10.2) because of the cache hierarchy.

Actually, buffer reads can use L1 and L2. When caching is not used for a buffer, reads from that buffer bypass L2. After a buffer read, the line is invalidated; then, on the next read, it is read again (from the same wavefront or from a different clause). After a buffer write, the changed parts of the cache line are written to memory.

Buffers and images are written through the texture L2 cache, but this is flushed immediately after an image write.

The data in private memory is first placed in registers. If more private memory is used than can be placed in registers, or dynamic indexing is used on private arrays, the overflow data is placed (spilled) into scratch memory. Scratch memory is a private subset of global memory, so performance can be dramatically degraded if spilling occurs.

Global memory can be in the high-speed GPU memory (VRAM) or in the host memory, which is accessed by the PCIe bus. A work-item can access global memory either as a buffer or a memory object. Buffer objects are generally read and written directly by the work-items. Data is accessed through the L2 and L1 data caches on the GPU. This limited form of caching provides read coalescing among work-items in a wavefront. Similarly, writes are executed through the texture L2 cache.

Global atomic operations are executed through the texture L2 cache. Atomic instructions that return a value to the kernel are handled similarly to fetch instructions: the kernel must use `S_WAITCNT` to ensure the results have been written to the destination GPR before using the data.

10.3 LDS Access

The LDS is accessed in one of three ways:

- Direct Read
- Parameter Read
- Indexed or Atomic

The following subsections describe these methods.

10.3.1 LDS Direct Reads

Direct reads are only available in LDS, not in GDS.

LDS Direct reads occur in vector ALU (VALU) instructions and allow the LDS to supply a single DWORD value which is broadcast to all threads in the wavefront

and is used as the SRC0 input to the ALU operations. A VALU instruction indicates that input is to be supplied by LDS by using the `LDS_DIRECT` for the SRC0 field.

The LDS address and data-type of the data to be read from LDS comes from the M0 register:

```
LDS_addr = M0[15:0]    (byte address and must be dword aligned)
DataType = M0[18:16]
    0 – unsigned byte
    1 – unsigned short
    2 – dword
    3 – unused
    4 – signed byte
    5 – signed short
```

10.3.2 LDS Parameter Reads

Parameter reads are only available in LDS, not in GDS.

Pixel shaders use LDS to read vertex parameter values; the pixel shader then interpolates them to find the per-pixel parameter values. LDS parameter reads occur when the following opcodes are used.

- `V_INTERP_P1_F32D = P10 * S + P0` Parameter interpolation, first step.
- `V_INTERP_P2_F32D = P20 * S + DP` Parameter interpolation, second step.
- `V_INTERP_MOV_F32D = {P10,P20,P0}[S]` Parameter load.

The typical parameter interpolation operations involves reading three parameters: P0, P10, and P20, and using the two barycentric coordinates, I and J, to determine the final per-pixel value:

$$\text{Final value} = P0 + P10 * I + P20 * J$$

Parameter interpolation instructions indicate the parameter attribute number (0 to 32) and the component number (0=x, 1=y, 2=z and 3=w).

Table 10.1 lists and briefly describes the parameter instruction fields.

Table 10.1 Parameter Instruction Fields

Field	Size	Description
VDST	8	Destination VGPR. Also acts as source for <code>v_interp_p2_f32</code> .
OP	2	Opcode: 0: <code>v_interp_p1_f32</code> $VDST = P10 * VSRC + P0$ 1: <code>v_interp_p2_f32</code> $VDST = P20 * VSRC + VDST$ 2: <code>v_interp_mov_f32</code> $VDST = (P0, P10 \text{ or } P20 \text{ selected by } VSRC[1:0])$ P0, P10 and P20 are parameter values read from LDS
ATTR	6	Attribute number: 0 to 32.
ATTR CHAN	2	0=X, 1=Y, 2=Z, 3=W
VSRC	8	Source VGPR supplies interpolation "I" or "J" value. For <code>OP==v_interp_mov_f32</code> : 0=P10, 1=P20, 2=P0. VSRC must not be the same register as VDST because 16-bank LDS chips implement <code>v_interp_p1</code> as a macro of two instructions.
(M0)	32	Use of the M0 register is automatic. M0 must contain: $\{ 1'b0, \text{new_prim_mask}[15:1], \text{lds_param_offset}[15:0] \}$

Parameter interpolation and parameter move instructions must initialize the M0 register before using it, as shown in Table 10.1. The `lds_param_offset[15:0]` is an address offset from the beginning of LDS storage allocated to this wavefront to where parameters begin in LDS memory for this wavefront. The `new_prim_mask` is a 15-bit mask with one bit per quad; a one in this mask indicates that this quad begins a new primitive, a zero indicates it uses the same primitive as the previous quad. The mask is 15 bits, not 16, since the first quad in a wavefront always begins a new primitive and so it is not included in the mask.

Figure 10.3 shows how parameters are laid out in LDS memory.

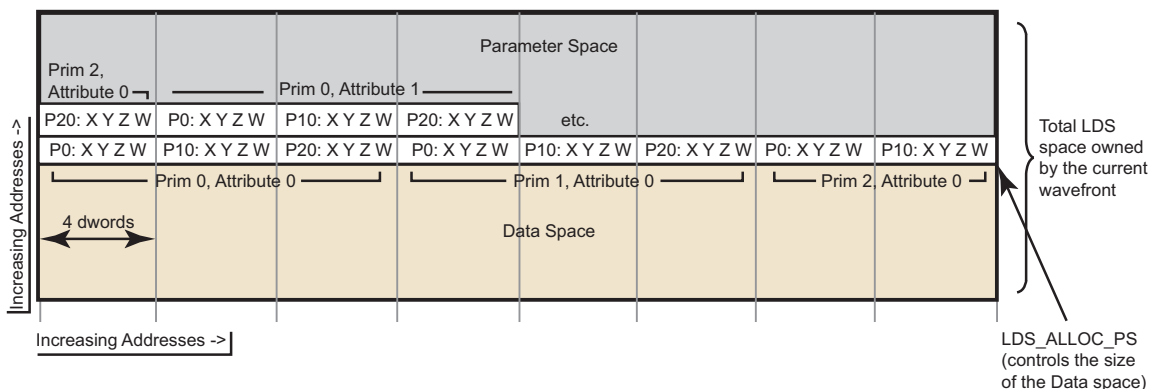


Figure 10.3 LDS Layout with Parameters and Data Share

10.3.3 Data Share Indexed and Atomic Access

Both LDS and GDS can perform indexed and atomic data share operations. For brevity, “LDS” is used in the text below and, except where noted, also applies to GDS.

Indexed and atomic operations supply a unique address per work-item from the VGPRs to the LDS, and supply or return unique data per work-item back to VGPRs. Due to the internal banked structure of LDS, operations can complete in as little as two cycles, or take as many 64 cycles, depending upon the number of bank conflicts (addresses that map to the same memory bank).

Indexed operations are simple LDS load and store operations that read data from, and return data to, VGPRs.

Atomic operations are arithmetic operations that combine data from VGPRs and data in LDS, and write the result back to LDS. Atomic operations have the option of returning the LDS “pre-op” value to VGPRs.

Table 10.2 lists and briefly describes the LDS instruction fields.

Table 10.2 LDS Instruction Fields

Field	Size	Description
OP	7	LDS opcode.
GDS	1	0 = LDS, 1 = GDS.
OFFSET0	8	Immediate offset, in bytes.
OFFSET1	8	Instructions with one address combine the offset fields into a single 16-bit unsigned offset: {offset1, offset0}. Instructions with two addresses (for example: READ2) use the offsets separately as two 8-bit unsigned offsets. DS_*_SRC2_* ops treat the offset as a 16-bit signed Dword offset.
VDST	8	VGPR to which result is written: either from LDS-load or atomic return value.
ADDR	8	VGPR that supplies the byte address offset.
DATA0	8	VGPR that supplies first data source.
DATA1	8	VGPR that supplies second data source.
(M0)	32	Implied use of M0. M0[16:0] contains the byte-size of the LDS segment. This is used to clamp the final address.

All LDS operations require that M0 be initialized prior to use. M0 contains a size value that can be used to restrict access to a subset of the allocated LDS range. If no clamping is wanted, set M0 to 0xFFFFFFFF.

Table 10.3 lists and describes the LDS indexed loads and stores.

Table 10.3 LDS Indexed Load/Store

DS_READ_{B32, B64, B96, B128, U8, I8, U16, I16}	Read one value per thread; sign extend to DWORD, if signed.
DS_READ2_{B32, B64}	Read two values at unique addresses.
DS_READ2ST64_{B32, B64}	Read 2 values at unique addresses, offset *= 64
DS_WRITE_{B32, B64, B96, B128, B8, B16}	Write one value.
DS_WRITE2_{B32, B64}	Write two values.
DS_WRITE2ST64_{B32, B64}	Write two values, offset *= 64.
DS_WRXCHG2_RTN_{B32, B64}	Exchange GPR with LDS-memory.
DS_WRXCHG2ST64_RTN_{B32, B64}	Exchange GPR with LDS-memory, offset *= 64.
DS_PERMUTE_B32	Forward permute. Does not write any LDS memory. LDS[dst] = src0 returnVal = LDS[thread_id] Where "thread_id" is 0..63.
DS_BPERMUTE_B32	Backward permute. Does not actually write any LDS memory. LDS[thread_id] = src0 Where "thread_id" is 0..63, returnVal = LDS[dst]

Single Address Instructions

$$\text{LDS_Addr} = \text{LDS_BASE} + \text{VGPR}[\text{ADDR}] + \{\text{InstrOffset1}, \text{InstrOffset0}\}$$

Double Address Instructions

$$\text{LDS_Addr0} = \text{LDS_BASE} + \text{VGPR}[\text{ADDR}] + \text{InstrOffset0}$$

$$\text{LDS_Addr1} = \text{LDS_BASE} + \text{VGPR}[\text{ADDR}] + \text{InstrOffset1}$$

Note that LDS_ADDR1 is used only for READ2*, WRITE2*, and WREXCHG2*.

M0[15:0] provides the size in bytes for this access. The size sent to LDS is MIN(M0, LDS_SIZE), where LDS_SIZE is the amount of LDS space allocated by the shader processor interpolator, SPI, at the time the wavefront was created.

The address comes from VGPR, and both ADDR and InstrOffset are byte addresses.

At the time of wavefront creation, LDS_BASE is assigned to the physical LDS region owned by this wavefront or work-group.

Specify only one address by setting both offsets to the same value. This causes only one read or write to occur and uses only the first DATA0.

LDS Atomic Ops

DS_<atomicOp> OP, GDS=0, OFFSET0, OFFSET1, VDST, ADDR, Data0, Data1

Data size is encoded in atomicOp: byte, word, Dword, or double.

$LDS_Addr0 = LDS_BASE + VGPR[ADDR] + \{InstrOffset1, InstrOffset0\}$

ADDR is a Dword address. VGPRs 0,1 and dst are double-GPRs for doubles data.

VGPR data sources can only be VGPRs or constant values, not SGPRs.

Chapter 11


Exporting Pixel Color and Vertex Shader Parameters

The export instruction copies pixel or vertex shader data from VGPRs into a dedicated output buffer. The export instruction outputs the following types of data.

- Vertex Position
- Vertex Parameter
- Pixel color
- Pixel depth (Z)

11.1 Microcode Encoding

The export instruction uses the EXP microcode format.



The fields are described in Table 11.1.

Table 11.1 EXP Encoding Field Descriptions

Field	Bits	Description
VM	12	Valid Mask. When set to 1, this indicates that the EXEC mask represents the valid-mask for this wavefront. It can be sent multiple times per shader (the final value is used), but must be sent at least once per pixel shader.
DONE	11	This is the final pixel shader or vertex-position export of the program. Used only for pixel and position exports. Set to zero for parameters.
COMPR	10	Compressed data. When set, indicates that the data being exported is 16-bits per component rather than the usual 32-bit.
TARGET	10:4	Indicates type of data exported. 0..7 MRT 0..7 8 Z 9 Null (no data) 12-15 Position 0..3 32-63 Param 0..31

Table 11.1 EXP Encoding Field Descriptions (Cont.)

Field	Bits	Description
EN	3:0	COMPR==1 : export half-dword enable. Valid values are: 0x0,3,C,F. [0] enables VSRC0 : R,G from one VGPR [2] enables VSRC1 : B,A from one VGPR COMPR==0 : [0-3] = enables for VSRC0..3. EN can be zero (used when exporting only valid mask to NULL target).
VSRC3	63:56	VGPR from which to read data. Pos & Param: vsrc0=X, 1=Y, 2=Z, 3=W MRT: vsrc0=R, 1=G, 2=B, 3=A
VSRC2	55:48	
VSRC1	47:40	
VSRC0	39:32	

11.2 Operations

11.2.1 Pixel Shader Exports

Export instructions copy color data to the MRTs. Data always has four components (R, G, B, A). Optionally, export instructions also output depth (Z) data.

Every pixel shader must have at least one export instruction. The last export instruction executed must have the DONE bit set to one.

The EXEC mask is applied to all exports. Only pixels with the corresponding EXEC bit set to 1 export data to the output buffer. Results from multiple exports are accumulated in the output buffer.

At least one export must have the VM bit set to 1. This export, in addition to copying data to the color or depth output buffer, also informs the color buffer which pixels are valid and which have been discarded. The value of the EXEC mask communicates the pixel valid mask. If multiple exports are sent with VM set to 1, the mask from the final export is used. If the shader program wants to only update the valid mask but not send any new data, the program can do an export to the NULL target.

11.2.2 Vertex Shader Exports

The vertex shader uses export instructions to output vertex position data and vertex parameter data to the output buffer. This data is passed on to subsequent pixel shaders.

Every vertex shader must output at least one position vector (x, y, z; w is optional) to the POS0 target. The last position export must have the DONE bit set to 1. A vertex shader can export zero or more parameters. For best performance, it is best to output all position data as early as possible in the vertex shader.

11.3 Dependency Checking

Export instructions are executed by the hardware in two phases. First, the instruction is selected to be executed, and EXPCNT is incremented by 1. At this time, the hardware requests the use of internal busses needed to complete the instruction.

When access to the bus is granted, the EXEC mask is read and the VGPR data sent out. After the last of the VGPR data is sent, the EXPCNT counter is decremented by 1.

Use S_WAITCNT on EXPCNT to prevent the shader program from overwriting EXEC or the VGPRs holding the data to be exported before the export operation has completed.

Multiple export instructions can be outstanding at one time. Exports of the same type (for example: position) are completed in order, but exports of different types can be completed out of order.

If the STATUS register's SKIP_EXPORT bit is set to one, the hardware treats all EXPORT instructions as if they were NOPs.

Chapter 12

Instruction Set

This chapter lists, and provides descriptions for, all instructions in the GCN Generation 3 environment. Instructions are grouped according to their format.

Instruction suffixes have the following definitions:

- B32 Bitfield (untyped data) 32-bit
- B64 Bitfield (untyped data) 64-bit
- F32 floating-point 32-bit (IEEE 754 single-precision float)
- F64 floating-point 64-bit (IEEE 754 double-precision float)
- I32 signed 32-bit integer
- I64 signed 64-bit integer
- U32 unsigned 32-bit integer
- U64 unsigned 64-bit integer

If an instruction has two suffixes (for example, `_I32_F32`), the first suffix indicates the destination type, the second the source type.

Note that `.u` or `.i` specifies to interpret the argument as an unsigned or signed float.

12.1 SOP2 Instructions

Instruction **S_ABSDIFF_I32**

Description $D.i = \text{abs}(S0.i - S1.i)$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 42 (0x2A)

1	0	OP	SDST	SSRC1	SSRC0	+0
---	---	----	------	-------	-------	----

Instruction **S_ADD_I32**

Description $D.u = S0.i + S1.i$. SCC = signed overflow.

Microcode SOP2 Opcode 2 (0x2)



Instruction **S_ADD_U32**

Description $D.u = S0.u + S1.u$. SCC = unsigned carry out.

Microcode SOP2 Opcode 0 (0x0)



Instruction **S_ADDC_U32**

Description $D.u = S0.u + S1.u + SCC$. SCC = unsigned carry-out.


Microcode SOP2 Opcode 4 (0x4)



Instruction **S_AND_B32**

Description $D.u = S0.u \& S1.u$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 13 (0xC)



Instruction **S_AND_B64**

Description $D.u = S0.u \& S1.u$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 13 (0xD)



Instruction **S_ANDN2_B32**

Description $D.u = S0.u \& \sim S1.u$. SCC = 1 if result is non-zero.


Microcode SOP2 Opcode 18 (0x12)



Instruction **S_ANDN2_B64**

Description $D.u = S0.u \& \sim S1.u$. SCC = 1 if result is non-zero.


Microcode SOP2 Opcode 19 (0x13)



Instruction **S_ASHR_I32**

Description $D.i = \text{signext}(S0.i) \gg S1.i[4:0]$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 32 (0x20)



Instruction **S_ASHR_I64**

Description $D.i = \text{signext}(S0.i) \gg S1.i[5:0]$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 33 (0x21)

1	0	OP	SDST	SSRC1	SSRC0	+0
---	---	----	------	-------	-------	----

Instruction **S_BFE_I32**

Description Replace description text with:
 DX11 Unsigned bitfield extract. Extract a contiguous range of bits from 32-bit source.
 SRC0 = input data
 SRC1 = the lowest bit position to select
 SRC2 = the width of the bit field
 Returns the bit starting at "offset" and ending at "offset+width-1".
 The final result is sign-extended.
 If (src2[4:0] == 0) {
 dst = 0;
 }
 Else if (src2[4:0] + src1[4:0] < 32) {
 dst = (src0 << (32-src1[4:0] - src2[4:0])) >>> (32 - src2[4:0])
 }
 Else {
 dst = src0 >>> src1[4:0]
 >>> means arithmetic shift right.
 }
 SCC = 1 if result is non-zero. Test sign-extended result.

Microcode SOP2 Opcode 38 (0x26)

1	0	OP	SDST	SSRC1	SSRC0	+0
---	---	----	------	-------	-------	----

Instruction **S_BFE_I64**

Description Bit field extract. S0 is data, S1[5:0] is field offset, S1[22:16] is field width.
 $D.i = (S0.u \gg S1.u[5:0]) \& ((1 \ll S1.u[22:16]) - 1)$. S
 CC = 1 if result is non-zero. Test sign-extended result.

Microcode SOP2 Opcode 40 (0x28)

1	0	OP	SDST	SSRC1	SSRC0	+0
---	---	----	------	-------	-------	----

Instruction **S_BFE_U32**

Description DX11 Unsigned bitfield extract. Extract a contiguous range of bits from 32-bit source.
 SRC0 = input data
 SRC1 = the lowest bit position to select
 SRC2 = the width of the bit field
 Returns the bit starting at "offset" and ending at "offset+width-1".

```

If (src2[4:0] == 0) {
    dst = 0;
}
Else if (src2[4:0] + src1[4:0] < 32) {
    dst = (src0 << (32-src1[4:0] - src2[4:0])) >> (32 - src2[4:0])
}
Else {
    dst = src0 >> src1[4:0]
}
SCC = 1 if result is non-zero. Test sign-extended result.
    
```

Microcode SOP2 Opcode 37 (0x25)



Instruction **S_BFE_U64**

Description Bit field extract. S0 is data, S1[4:0] is field offset, S1[22:16] is field width.
 D.u = (S0.u >> S1.u[5:0]) & ((1 << S1.u[22:16]) - 1). SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 39 (0x27)



Instruction **S_BFM_B32**

Description D.u = ((1 << S0.u[4:0]) - 1) << S1.u[4:0]; bitfield mask.

Microcode SOP2 Opcode 34 (0x22)



Instruction **S_BFM_B64**

Description $D.u = ((1 \ll S0.u[5:0]) - 1) \ll S1.u[5:0]$; bitfield mask.

Microcode SOP2 Opcode 35 (0x23)

1	0	OP	SDST	SSRC1	SSRC0	+0
---	---	----	------	-------	-------	----

Instruction **S_CBRANCH_G_FORK**

Description Conditional branch using branch stack. Arg0 = compare mask (VCC or any SGPR), Arg1 = 64-bit byte address of target instruction. See Section 4.6, on page 4-5.

Microcode SOP2 Opcode 41 (0x29)

1	0	OP	SDST	SSRC1	SSRC0	+0
---	---	----	------	-------	-------	----

Instruction **S_CSELECT_B32**

Description $D.u = SCC ? S0.u : S1.u$.

Microcode SOP2 Opcode 10 (0xA)

1	0	OP	SDST	SSRC1	SSRC0	+0
---	---	----	------	-------	-------	----

Instruction **S_CSELECT_B64**

Description $D.u = SCC ? S0.u : S1.u$.

Microcode SOP2 Opcode 11 (0xB)

1	0	OP	SDST	SSRC1	SSRC0	+0
---	---	----	------	-------	-------	----

Instruction **S_LSHL_B32**

Description $D.u = S0.u \ll S1.u[4:0]$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 28 (0x1C)

1	0	OP	SDST	SSRC1	SSRC0	+0
---	---	----	------	-------	-------	----

Instruction **S_LSHL_B64**

Description $D.u = S0.u \ll S1.u[5:0]$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 29 (0x1D)

1	0	OP	SDST	SSRC1	SSRC0	+0
---	---	----	------	-------	-------	----

Instruction **S_LSHR_B32**

Description $D.u = S0.u \gg S1.u[4:0]$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 30 (0x1E)

1	0	OP	SDST	SSRC1	SSRC0	+0
---	---	----	------	-------	-------	----

Instruction **S_LSHR_B64**

Description $D.u = S0.u \gg S1.u[5:0]$. SCC = 1 if result is non-zero.


Microcode SOP2 Opcode 31 (0x15)

1	0	OP	SDST	SSRC1	SSRC0	+0
---	---	----	------	-------	-------	----

Instruction **S_MAX_I32**

Description $D.i = (S0.i > S1.i) ? S0.i : S1.i$. SCC = 1 if S0 is max.


Microcode SOP2 Opcode 8 (0x8)



Instruction **S_MAX_U32**

Description $D.u = (S0.u > S1.u) ? S0.u : S1.u$. SCC = 1 if S0 is max.


Microcode SOP2 Opcode 9 (0x9)



Instruction **S_MIN_I32**

Description $D.i = (S0.i < S1.i) ? S0.i : S1.i$. SCC = 1 if S0 is min.


Microcode SOP2 Opcode 6 (0x6)



Instruction **S_MIN_U32**

Description $D.u = (S0.u < S1.u) ? S0.u : S1.u$. SCC = 1 if S0 is min.


Microcode SOP2 Opcode 7 (0x7)



Instruction **S_MUL_I32**

Description $D.i = S0.i * S1.i.$


Microcode SOP2 Opcode 36 (0x24)



Instruction **S_NAND_B32**

Description $D.u = \sim(S0.u \& S1.u).$ SCC = 1 if result is non-zero.


Microcode SOP2 Opcode 22 (0x16)



Instruction **S_NAND_B64**

Description $D.u = \sim(S0.u \& S1.u).$ SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 23 (0x17)



Instruction **S_NOR_B32**

Description $D.u = \sim(S0.u | S1.u).$ SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 24 (0x18)



Instruction **S_NOR_B64**

Description $D.u = \sim(S0.u \mid S1.u)$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 25 (0x19)



Instruction **S_OR_B32**

Description $D.u = S0.u \mid S1.u$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 14 (0xE)



Instruction **S_OR_B64**

Description $D.u = S0.u \mid S1.u$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 15 (0xF)



Instruction **S_ORN2_B32**

Description $D.u = S0.u \mid \sim S1.u$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 20 (0x 14)



Instruction **S_ORN2_B64**

Description $D.u = S0.u \mid \sim S1.u$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 21 (0x15)



Instruction **S_SUB_I32**

Description $D.u = S0.i - S1.i$. SCC = borrow.


Microcode SOP2 Opcode 3 (0x3)



Instruction **S_SUB_U32**

Description $D.u = S0.u - S1.u$. SCC = unsigned carry out.


Microcode SOP2 Opcode 1 (0x1)



Instruction **S_SUBB_U32**

Description $D.u = S0.u - S1.u - SCC$. SCC = unsigned carry-out.

Microcode SOP2 Opcode 5 (0x5)



Instruction **S_XNOR_B32**

Description $D.u = \sim(S0.u \wedge S1.u)$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 26 (0x1A)



Instruction **S_XNOR_B64**

Description $D.u = \sim(S0.u \wedge S1.u)$. SCC = 1 if result is non-zero.


Microcode SOP2 Opcode 27 (0x1B)



Instruction **S_XOR_B32**

Description $D.u = S0.u \wedge S1.u$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 16 (0x10)



Instruction **S_XOR_B64**

Description $D.u = S0.u \wedge S1.u$. SCC = 1 if result is non-zero.

Microcode SOP2 Opcode 17 (0x11)



12.2 SOPK Instructions

Instruction **S_ADDK_I32**

Description $D.i = D.i + \text{signext}(\text{SIMM16})$. SCC = signed overflow.


Microcode SOPK Opcode 14 (0xE)



Instruction **S_CBRANCH_I_FORK**

Description Conditional branch using branch-stack. Arg0(sdst) = compare mask (VCC or any SGPR), SIMM16 = signed DWORD branch offset relative to next instruction. See Section 4.6, on page 4-5.


Microcode SOPK Opcode 16 (0x10)



Instruction **S_CMOVK_I32**

Description if (SCC) $D.i = \text{signext}(\text{SIMM16})$; else NOP.


Microcode SOPK Opcode 1 (0x1)



Instruction **S_CMPK_EQ_I32**

Description $\text{SCC} = (D.i == \text{signext}(\text{SIMM16}))$.

Microcode SOPK Opcode 2 (0x2)



Instruction **S_CMPK_EQ_U32**

Description $SCC = (D.u == SIMM16).$

Microcode SOPK Opcode 8 (0x8)



Instruction **S_CMPK_GE_I32**

Description $SCC = (D.i \geq \text{signext}(SIMM16)).$

Microcode SOPK Opcode 5 (0x5)



Instruction **S_CMPK_GE_U32**

Description $SCC = (D.u \geq SIMM16).$


Microcode SOPK Opcode 11 (0xB)



Instruction **S_CMPK_GT_I32**

Description $SCC = (D.i > \text{signext}(SIMM16)).$

Microcode SOPK Opcode 4 (0x4)



Instruction **S_CMPK_GT_U32**

Description $SCC = (D.u > SIMM16).$

Microcode SOPK Opcode 10 (0xA)



Instruction **S_CMPK_LE_I32**

Description $SCC = (D.i \leq \text{signext}(SIMM16)).$


Microcode SOPK Opcode 7 (0x7)



Instruction **S_CMPK_LE_U32**

Description $D.u = SCC = (D.u \leq SIMM16).$

Microcode SOPK Opcode 13 (0xD)



Instruction **S_CMPK_LG_I32**

Description $SCC = (D.i \neq \text{signext}(SIMM16)).$

Microcode SOPK Opcode 3 (0x3)



Instruction **S_CMPK_LG_U32**

Description $SCC = (D.u \neq SIMM16).$

Microcode SOPK Opcode 9 (0x9)



Instruction **S_CMPK_LT_I32**

Description $SCC = (D.i < \text{signext}(SIMM16)).$


Microcode SOPK Opcode 6 (0x6)



Instruction **S_CMPK_LT_U32**

Description $SCC = (D.u < SIMM16).$

Microcode SOPK Opcode 12 (0xC)



Instruction **S_GETREG_B32**

Description D.u = hardware register. Read some or all of a hardware register into the LSBs of D. See Table 5.7 on page 5-8. $SIMM16 = \{\text{size}[4:0], \text{offset}[4:0], \text{hwRegId}[5:0]\}$; offset is in the range from 0 to 31, size is in the range from 1 to 32.

Microcode SOPK Opcode 17 (0x11)



Instruction **S_MOVK_I32**

Description $D.i = \text{signext}(\text{SIMM16})$.


Microcode SOPK Opcode 0 (0x0)



Instruction **S_MULK_I32**

Description $D.i = D.i * \text{signext}(\text{SIMM16})$. SCC = overflow.

Microcode SOPK Opcode 15 (0xF)



Instruction **S_SETREG_B32**

Description Hardware register = D.u. Write some or all of the LSBs of D into a hardware register (note that D is a source SGPR). See Table 5.7 on page 5-8.
 $\text{SIMM16} = \{\text{size}[4:0], \text{offset}[4:0], \text{hwRegId}[5:0]\}$; offset is in the range from 0 to 31, size is in the range from 1 to 32.


Microcode SOPK Opcode 18 (0x12)



Instruction **S_SETREG_IMM32_B32**

Description This instruction uses a 32-bit literal constant. Write some or all of the LSBs of SIMM32 into a hardware register.
SIMM16 = {size[4:0], offset[4:0], hwRegId[5:0]}; offset is 0–31, size is 1–32.

Microcode SOPK Opcode 20 (0x14)



12.3 SOP1 Instructions

Instruction **S_ABS_I32**

Description $D.i = \text{abs}(S0.i)$. SCC=1 if result is non-zero.

Microcode SOP1 Opcode 48 (0x30)



Instruction **S_AND_SAVEEXEC_B64**

Description $D.u = \text{EXEC}$, $\text{EXEC} = S0.u \ \& \ \text{EXEC}$. SCC = 1 if the new value of EXEC is non-zero.

Microcode SOP1 Opcode 32 (0x20)



Instruction **S_ANDN2_SAVEEXEC_B64**

Description $D.u = \text{EXEC}$, $\text{EXEC} = S0.u \ \& \ \sim\text{EXEC}$. SCC = 1 if the new value of EXEC is non-zero.

Microcode SOP1 Opcode 35 (0x23)



Instruction **S_BCNT0_I32_B32**

Description $D.i = \text{CountZeroBits}(S0.u)$. SCC = 1 if result is non-zero.

Microcode SOP1 Opcode 10 (0xA)



Instruction **S_BCNT0_I32_B64**

Description $D.i = \text{CountZeroBits}(S0.u)$. SCC = 1 if result is non-zero.

Microcode SOP1 Opcode 11 (0xB)



Instruction **S_BCNT1_I32_B32**

Description $D.i = \text{CountOneBits}(S0.u)$. SCC = 1 if result is non-zero.

Microcode SOP1 Opcode 12 (0xC)



Instruction **S_BCNT1_I32_B64**

Description $D.i = \text{CountOneBits}(S0.u)$. SCC = 1 if result is non-zero.

Microcode SOP1 Opcode 13 (0xD)



Instruction **S_BITSET0_B32**

Description $D.u[S0.u[4:0]] = 0$.

Microcode SOP1 Opcode 24 (0x18)



Instruction **S_BITSET0_B64**

Description $D.u[S0.u[5:0]] = 0.$


Microcode SOP1 Opcode 25 (0x19)



Instruction **S_BITSET1_B32**

Description $D.u[S0.u[4:0]] = 1.$

Microcode SOP1 Opcode 26 (0x1A)



Instruction **S_BITSET1_B64**

Description $D.u[S0.u[5:0]] = 1.$


Microcode SOP1 Opcode 27 (0x1B)



Instruction **S_BREV_B32**

Description $D.u = S0.u[0:31]$ (reverse bits).

Microcode SOP1 Opcode 8 (0x8)



Instruction **S_BREV_B64**

Description D.u = S0.u[0:63] (reverse bits).

Microcode SOP1 Opcode 9 (0x9)



Instruction **S_CBRANCH_JOIN**

Description Conditional branch join point. Arg0 = saved CSP value. No dest. See Section 4.6, on page 4-5.


Microcode SOP1 Opcode 46 (0x2E)



Instruction **S_CMOV_B32**

Description if(SCC) D.u = S0.u; else NOR.

Microcode SOP1 Opcode 2 (0x2)



Instruction **S_CMOV_B64**

Description if(SCC) D.u = S0.u; else NOR.

Microcode SOP1 Opcode 3 (0x3)



Instruction **S_FF0_I32_B32**

Description D.i = FindFirstZero(S0.u) from LSB; if no zeros, return -1.


Microcode SOP1 Opcode 14 (0xE)



Instruction **S_FF0_I32_B64**

Description D.i = FindFirstZero(S0.u) from LSB; if no zeros, return -1.


Microcode SOP1 Opcode 15 (0xF)



Instruction **S_FF1_I32_B32**

Description D.i = FindFirstOne(S0.u) from LSB; if no ones, return -1.


Microcode SOP1 Opcode 16 (0x10)



Instruction **S_FF1_I32_B64**

Description D.i = FindFirstOne(S0.u) from LSB; if no ones, return -1.

Microcode SOP1 Opcode 17 (0x11)



Instruction **S_FLBIT_I32**

Description D.i = Find first bit opposite of sign bit from MSB. If S0 == -1, return -1.

Microcode SOP1 Opcode 20 (0x14)



Instruction **S_FLBIT_I32_B32**

Description D.i = FindFirstOne(S0.u) from MSB; if no ones, return -1.

Microcode SOP1 Opcode 18 (0x12)



Instruction **S_FLBIT_I32_B64**

Description D.i = FindFirstOne(S0.u) from MSB; if no ones, return -1.

Microcode SOP1 Opcode 19 (0x13)



Instruction **S_FLBIT_I32_I64**

Description D.i = Find first bit opposite of sign bit from MSB. If S0 == -1, return -1.

Microcode SOP1 Opcode 21 (0x15)



Instruction **S_GETPC_B64**

Description $D.u = PC + 4$; destination receives the byte address of the next instruction.


Microcode SOP1 Opcode 28 (0x1C)



Instruction **S_MOV_B32**

Description $D.u = S0.u$.


Microcode SOP1 Opcode 1 (0x0)



Instruction **S_MOV_B64**

Description $Du = S0.u$.


Microcode SOP1 Opcode 1 (0x1)



Instruction **S_MOVRELD_B32**

Description $SGPR[D.u + M0.u] = SGPR[S0.u]$.

Microcode SOP1 Opcode 44 (0x2C)



Instruction **S_MOVRELD_B64**

Description $SGPR[D.u + M0.u] = SGPR[S0.u]$. M0 and D.u must be even.

Microcode SOP1 Opcode 45 (0x2D)



Instruction **S_MOVRELS_B32**

Description $SGPR[D.u] = SGPR[S0.u + M0.u]$.

Microcode SOP1 Opcode 42 (0x2A)



Instruction **S_MOVRELS_B64**

Description $SGPR[D.u] = SGPR[S0.u + M0.u]$. M0 and S0.u must be even.

Microcode SOP1 Opcode 43 (0x2B)



Instruction **S_NAND_SAVEEXEC_B64**

Description $D.u = EXEC$, $EXEC = \sim(S0.u \& EXEC)$. SCC = 1 if the new value of EXEC is non-zero.

Microcode SOP1 Opcode 37 (0x25)



Instruction **S_NOR_SAVEEXEC_B64**

Description $D.u = EXEC, EXEC = \sim(S0.u \mid EXEC)$. SCC = 1 if the new value of EXEC is non-zero.


Microcode SOP1 Opcode 38 (0x26)



Instruction **S_NOT_B32**

Description $D.u = \sim S0.u$. SCC = 1 if result non-zero.


Microcode SOP1 Opcode 4 (0x4)



Instruction **S_NOT_B64**

Description $D.u = \sim S0.u$. SCC = 1 if result non-zero.


Microcode SOP1 Opcode 5 (0x5)



Instruction **S_OR_SAVEEXEC_B64**

Description $D.u = EXEC, EXEC = S0.u \mid EXEC$. SCC = 1 if the new value of EXEC is non-zero.

Microcode SOP1 Opcode 33 (0x21)



Instruction **S_ORN2_SAVEEXEC_B64**

Description $D.u = EXEC, EXEC = S0.u \mid \sim EXEC$. SCC = 1 if the new value of EXEC is non-zero.

Microcode SOP1 Opcode 36 (0x24)



Instruction **S_QUADMASK_B32**

Description $D.u = QuadMask(S0.u)$. $D[0] = OR(S0[3:0])$, $D[1] = OR(S0[7:4])$. SCC = 1 if result is non-zero.

Microcode SOP1 Opcode 40 (0x28)



Instruction **S_QUADMASK_B64**

Description $D.u = QuadMask(S0.u)$. $D[0] = OR(S0[3:0])$, $D[1] = OR(S0[7:4])$. SCC = 1 if result is non-zero.

Microcode SOP1 Opcode 41 (0x29)



Instruction **S_RFE_B64**

Description Return from Exception; $PC = S0.u$. This instruction sets PRIV to 0.

Microcode SOP1 Opcode 31 (0x1F)



Instruction **S_SET_GPR_IDX_IDX**

Description $M0[7:0] = S0.U[7:0]$. Modify the index used in vector GPR indexing.

Microcode SOP1 Opcode 50 (0x32)



Instruction **S_SETPC_B64**

Description $PC = S0.u$; $S0.u$ is a byte address of the instruction to jump to.


Microcode SOP1 Opcode 29 (0x1D)



Instruction **S_SEXT_I32_I8**

Description $D.i = \text{signext}(S0.i[7:0])$.


Microcode SOP1 Opcode 22 (0x16)



Instruction **S_SEXT_I32_I16**

Description $D.i = \text{signext}(S0.i[15:0])$.

Microcode SOP1 Opcode 23 (0x17)



Instruction **S_SWAPPC_B64**

Description $D.u = PC + 4; PC = S0.u.$

Microcode SOP1 Opcode 30 (0x1E)

1	0	1	1	1	1	1	1	0	1	SDST	OP	SSRC0	+0
---	---	---	---	---	---	---	---	---	---	------	----	-------	----

Instruction **S_WQM_B32**

Description $D.u = \text{WholeQuadMode}(S0.u).$ SCC = 1 if result is non-zero.
 Apply whole quad mode to the bitmask specified in SSRC0. Whole quad mode checks each group of four bits in the bitmask; if any bit is set to 1, all four bits are set to 1 in the result. This operation is repeated for the entire bitmask.

Microcode SOP1 Opcode 6 (0x6)

1	0	1	1	1	1	1	1	0	1	SDST	OP	SSRC0	+0
---	---	---	---	---	---	---	---	---	---	------	----	-------	----

Instruction **S_WQM_B64**

Description $D.u = \text{WholeQuadMode}(S0.u).$ SCC = 1 if result is non-zero.
 Apply whole quad mode to the bitmask specified in SSRC0. Whole quad mode checks each group of four bits in the bitmask; if any bit is set to 1, all four bits are set to 1 in the result. This operation is repeated for the entire bitmask.

Microcode SOP1 Opcode 7 (0x7)

1	0	1	1	1	1	1	1	0	1	SDST	OP	SSRC0	+0
---	---	---	---	---	---	---	---	---	---	------	----	-------	----

Instruction **S_XNOR_SAVEEXEC_B64**

Description $D.u = \text{EXEC}, \text{EXEC} = \sim(S0.u \wedge \text{EXEC}).$ SCC = 1 if the new value of EXEC is non-zero.


Microcode SOP1 Opcode 39 (0x27)

1	0	1	1	1	1	1	1	0	1	SDST	OP	SSRC0	+0
---	---	---	---	---	---	---	---	---	---	------	----	-------	----

Instruction **S_XOR_SAVEEXEC_B64**

Description $D.u = EXEC, EXEC = S0.u \wedge EXEC$. SCC = 1 if the new value of EXEC is non-zero.

Microcode SOP1 Opcode 34 (0x22)



12.4 SOPC Instructions

Instruction **S_BITCMP0_B32**

Description $SCC = (S0.u[S1.u[4:0]] == 0).$

Microcode SOPC Opcode 12 (C)



Instruction **S_BITCMP0_B64**

Description $SCC = (S0.u[S1.u[5:0]] == 0).$


Microcode SOPC Opcode 14 (0xE)



Instruction **S_BITCMP1_B32**

Description $SCC = (S0.u[S1.u[4:0]] == 1).$

Microcode SOPC Opcode 13 (0xD)



Instruction **S_BITCMP1_B64**

Description $SCC = (S0.u[S1.u[5:0]] == 1).$

Microcode SOPC Opcode 15 (0xF)



Instruction **S_CMP_EQ_I32**

Description $SCC = (S0.i == S1.i).$

Microcode SOPC Opcode 0 (0x0)



Instruction **S_CMP_EQ_U32**

Description $SCC = (S0.u == S1.u).$

Microcode SOPC Opcode 6 (0x6)



Instruction **S_CMP_EQ_U64**

Description $SCC = (S0.i64 == S1.i64).$

Microcode SOPC Opcode 18 (0x12)



Instruction **S_CMP_GE_I32**

Description $SCC = (S0.i >= S1.i).$

Microcode SOPC Opcode 3 (0x3)



Instruction **S_CMP_GE_U32**

Description $SCC = (S0.u \geq S1.u).$

Microcode SOPC Opcode 9 (0x9)



Instruction **S_CMP_GT_I32**

Description $SCC = (S0.i > S1.i).$


Microcode SOPC Opcode 2 (0x2)



Instruction **S_CMP_GT_U32**

Description $SCC = (S0.u > S1.u).$

Microcode SOPC Opcode 8 (0x8)



Instruction **S_CMP_LE_I32**

Description $SCC = (S0.i \leq S1.i).$

Microcode SOPC Opcode 5 (0x5)



Instruction **S_CMP_LE_U32**

Description $SCC = (S0.u \leq S1.u).$

Microcode SOPC Opcode 11 (0xB)



Instruction **S_CMP_LG_I32**

Description $SCC = (S0.i \neq S1.i).$


Microcode SOPC Opcode 1 (0x1)



Instruction **S_CMP_LG_U32**

Description $SCC = (S0.u \neq S1.u).$

Microcode SOPC Opcode 7 (0x7)



Instruction **S_CMP_LG_U64**

Description $SCC = (S0.i64 \neq S1.i64).$


Microcode SOPC Opcode 19 (0x13)



Instruction **S_CMP_LT_I32**

Description $SCC = (S0.i < S1.i).$

Microcode SOPC Opcode 4 (0x4)



Instruction **S_CMP_LT_U32**

Description $SCC = (S0.u < S1.u).$


Microcode SOPC Opcode 10 (0xA)



Instruction **S_CMP_NE_U64**

Description $SXCCX = (S0 \neq S1).$

Microcode SOPC Opcode 10 (0xA)



Instruction **S_SET_GPR_IDX_ON**

Description Enable GPR indexing mode. Vector operations after this perform relative GPR addressing based on the contents of M0. The structure SQ_M0_GPR_IDX_WORD can be used to decode M0. The raw contents of the S1 field are read and used to set the enable bits. S1[0] = VSRC0_REL, S1[1] = VSRC1_REL, S1[2] = VSRC2_REL, and S1[3] = VDST_REL.
 MODE.gpr_idx_en = 1;
 M0[7:0] = S0.u[7:0];
 M0[15:12] = SIMM4 (direct contents of S1 field);
 Remaining bits of M0 are unmodified.

Microcode SOPC Opcode 17 (0x11)

1	0	1	1	1	1	1	1	1	0	OP	SSRC1	SSRC0	+0
---	---	---	---	---	---	---	---	---	---	----	-------	-------	----

Instruction **S_SETVSKIP**

Description VSKIP = S0.u[S1.u[4:0]].
 Extract one bit from the SSRC0 SGPR, and use that bit to enable or disable VSKIP mode. In some cases, VSKIP mode can be used to skip over sections of code more quickly than branching. When VSKIP is enabled, the following instruction types are not executed: Vector ALU, Vector Memory, LDS, GDS, and Export.

Microcode SOPC Opcode 16 (0x10)

1	0	1	1	1	1	1	1	1	0	OP	SSRC1	SSRC0	+0
---	---	---	---	---	---	---	---	---	---	----	-------	-------	----

12.5 SOPP Instructions

Instruction **S_BARRIER**

Description Sync waves within a work-group.


Microcode SOPP Opcode 10 (0xA)



Instruction **S_BRANCH**

Description $PC = PC + \text{signext}(\text{SIMM16} * 4) + 4.$


Microcode SOPP Opcode 2 (0x2)



Instruction **S_CBRANCH_CDBGSYS**

Description Conditional branch when the SYStem debug bit is set.
 $\text{if}(\text{conditional_debug_system} \neq 0) \text{ then } PC = PC + \text{signext}(\text{SIMM16} * 4) + 4;$
 else NOP.

Microcode SOPP Opcode 23 (0x17)



Instruction **S_CBRANCH_CDBGSYS_AND_USER**

Description Conditional branch when both the SYStem and USER debug bits are set.
 $\text{if}(\text{conditional_debug_system} \ \&\& \ \text{conditional_debug_user}) \text{ then } PC = PC + \text{signext}(\text{SIMM16} * 4) + 4;$
 else NOP.

Microcode SOPP Opcode 26 (0x1A)



Instruction **S_CBRANCH_CDBGSYS_OR_USER**

Description Conditional branch when either the SYStem or USER debug bits are set.
 if(conditional_debug_system || conditional_debug_user) then PC = PC + signext(SIMM16 * 4) + 4; else NOP.

Microcode SOPP Opcode 25 (0x19)



Instruction **S_CBRANCH_CDBGUSER**

Description Conditional branch when the USER debug bit is set.
 if(conditional_debug_user != 0) then PC = PC + signext(SIMM16 * 4) + 4;
 else NOP.


Microcode SOPP Opcode 24 (0x18)



Instruction **S_CBRANCH_EXECSZ**

Description if(EXEC != 0) then PC = PC + signext(SIMM16 * 4) + 4; else NOP.


Microcode SOPP Opcode 9 (0x9)



Instruction **S_CBRANCH_EXECSZ**

Description if(EXEC == 0) then PC = PC + signext(SIMM16 * 4) + 4; else NOP.


Microcode SOPP Opcode 8 (0x8)



Instruction **S_CBRANCH_SCC0**

Description if(SCC == 0) then PC = PC + signext(SIMM16 * 4) + 4; else NOP


Microcode SOPP Opcode 4 (0x4)



Instruction **S_CBRANCH_SCC1**

Description if(SCC == 1) then PC = PC + signext(SIMM16 * 4) + 4; else NOP


Microcode SOPP Opcode 5 (0x5)



Instruction **S_CBRANCH_VCCNZ**

Description if(VCC != 0) then PC = PC + signext(SIMM16 * 4) + 4; else NOP


Microcode SOPP Opcode 7 (0x7)



Instruction **S_CBRANCH_VCCZ**

Description if(VCC == 0) then PC = PC + signext(SIMM16 * 4) + 4; else NOP

Microcode SOPP Opcode 6 (0x6)



Instruction **S_DECPERFLEVEL**

Description Decrement performance counter specified in SIMM16[3:0] by 1.

Microcode SOPP Opcode 21 (0x15)



Instruction **S_ENDPGM**

Description End of program; terminate wavefront.


Microcode SOPP Opcode 1 (0x1)



Instruction **S_ENDPGM_SAVED**

Description End of program; signal that a wave has been saved by the context-switch trap handler, and terminate wavefront. The hardware implicitly executes S_WAITCNT 0 before executing this instruction. Use S_ENDPGM in all cases unless you are executing the context-switch save handler.


Microcode SOPP Opcode 27 (0x1B)



Instruction **S_ICACHE_INV**

Description Invalidate entire L1 instruction cache.

Microcode SOPP Opcode 19 (0x13)



Instruction **S_INCPERFLEVEL**

Description Increment performance counter specified in SIMM16[3:0] by 1.

Microcode SOPP Opcode 20 (0x14)



Instruction **S_NOP**

Description Do nothing. Repeat NOP 1..8 times based on SIMM16[2:0]. 0 = 1 time, 7 = 8 times.

Microcode SOPP Opcode 0 (0x0)




Instruction **S_SENDMSG**

Description Send a message.

SIMM[3:0]	Message	Payload
1	interrupt	M0[7:0] carries user data. IDs are also sent (wave_id, cu_id, ...).
2	Gs	SIMM[5:4] defines GS_OP.
3	Gs_done	
4-14	unused	
15	System	Hardware internal use only.
SIMM[5:4]	GS OP	Payload
0	NOP	Use for gs-done only. M0[7:0] = gs-waveID
1	cut	SIMM[9:8] = stream_id EXEC is also sent. M0[7:0] = gs-waveID
2	emit	
3	emit-cut	


Microcode SOPP Opcode 16 (0x10)



Instruction **S_SENDMSGHALT**

Description Send a message and then HALT.

Microcode SOPP Opcode 17 (0x11)



Instruction **S_SET_GPR_IDX_MODE**

Description M0[15:12] = SIMM4. Modify the mode used for vector GPR indexing. The raw contents of the source field are read and used to set the enable bits. SIMM4[0] = VSRC0_REL, SIMM4[1] = VSRC1_REL, SIMM4[2] = VSRC2_REL and SIMM4[3] = VDST_REL.


Microcode SOPP Opcode 29 (0x1D)



Instruction **S_SET_GPR_IDX_OFF**

Description MODE.gpr_idx_en = 0. Clear GPR indexing mode. Vector operations after this do not perform relative GPR addressing regardless of the contents of M0. This instruction does not modify M0.


Microcode SOPP Opcode 28 (0x1C)



Instruction **S_SETHALT**

Description set HALT bit to value of SIMM16[0]. 1=halt, 0=resume. Halt is ignored while priv=1.


Microcode SOPP Opcode 13 (0xD)



Instruction **S_SETKILL**

Description Set KILL bit to value of SIMM16[0].

Microcode SOPP Opcode 11 (0xB)



Instruction **S_SETPRIO**

Description User-settable wave priority. The priority value is indicated in the two LSBs of the SIMM field.
0 = lowest, 3 = highest.


Microcode SOPP Opcode 15 (0xF)



Instruction **S_SLEEP**

Description Cause a wave to sleep for approximately 64*SIMM16[2:0] clocks.

Microcode SOPP Opcode 14 (0xE)



Instruction **S_TRAP**

Description Enter the trap handler. TrapID = SIMM16[7:0]. Wait for all instructions to complete, save {pc_rewind,trapID,pc} into tmp0,1; load TBA into PC, set PRIV=1 and continue. A trapID of zero is not allowed.


Microcode SOPP Opcode 18 (0x12)



Instruction **S_TTRACEDATA**

Description Send M0 as user data to thread-trace.


Microcode SOPP Opcode 22 (0x16)



Instruction **S_WAITCNT**

Description Wait for count of outstanding lds, vector-memory and export/vmem-write-data to be at or below the specified levels. simm16[3:0] = vmcount, simm16[6:4] = export/mem-write-data count, simm16[12:8] = LGKM_cnt (scalar-mem/GDS/LDS count). See Section 4.4, on page 4-2.

Microcode SOPP Opcode 12 (0xC)




12.6 SMEM Instructions

Instruction **S_ATC_PROBE**

Description Probe or prefetch an address into the SQC data cache.


Microcode SMEM Opcode 38 (0x26)



Instruction **S_ATC_PROBE_BUFFER**

Description Probe or prefetch an address into the SQC data cache. This instruction is used to probe buffers.

Microcode SMEM Opcode 38 (0x26)




Instruction **S_BUFFER_LOAD_DWORD**

Description Read one Dword from read-only memory describe by a buffer a constant (v#) through the constant cache (kcache).

```

m_offset      = IMM ? OFFSET : SGPR[OFFSET]
m_base       = { SGPR[SBASE * 2 +1][15:0], SGPR[SBASE] }
m_stride     = SGPR[SBASE * 2 +1][31:16]
m_num_records = SGPR[SBASE * 2 + 2]
m_size       = (m_stride == 0) ? 1 : m_num_records
m_addr       = (SGPR[SBASE * 2] + m_offset) & ~0x3
SGPR[SDST]   = read_dword_from_kcache(m_base, m_offset, m_size)
    
```

Microcode SMEM Opcode 8 (0x8)




Instruction **S_BUFFER_LOAD_DWORDX2**

Description Read two Dwords from read-only memory describe by a buffer a constant (v#) through the constant cache (kcache).

```

m_offset = IMM ? OFFSET : SGPR[OFFSET]
m_base = { SGPR[SBASE * 2 +1][15:0], SGPR[SBASE * 2] }
m_stride = SGPR[SBASE * 2 +1][31:16]
m_num_records = SGPR[SBASE * 2 + 2]
m_size = (m_stride == 0) ? 1 : m_num_records
m_addr = (SGPR[SBASE * 2] + m_offset) & ~0x3
SGPR[SDST] = read_dword_from_kcache(m_base, m_offset, m_size)
SGPR[SDST + 1] = read_dword_from_kcache(m_base, m_offset + 4, m_size)
    
```

Microcode SMEM Opcode 9 (0x9)




Instruction **S_BUFFER_LOAD_DWORDX4**

Description Read four Dwords from read-only memory describe by a buffer a constant (v#) through the constant cache (kcache).

```

m_offset = IMM ? OFFSET : SGPR[OFFSET]
m_base = { SGPR[SBASE * 2 + 1][15:0], SGPR[SBASE * 2] }
m_stride = SGPR[SBASE * 2 + 1][31:16]
m_num_records = SGPR[SBASE * 2 + 2]
m_size = (m_stride == 0) ? 1 : m_num_records
m_addr = (SGPR[SBASE * 2] + m_offset) & ~0x3
SGPR[SDST] = read_dword_from_kcache(m_base, m_offset, m_size)
SGPR[SDST + 1] = read_dword_from_kcache(m_base, m_offset + 4, m_size)
SGPR[SDST + 2] = read_dword_from_kcache(m_base, m_offset + 8, m_size)
SGPR[SDST + 3] = read_dword_from_kcache(m_base, m_offset + 12, m_size)
    
```

Microcode SMEM Opcode 10 (0xA)




Instruction **S_BUFFER_LOAD_DWORDX8**

Description Read eight Dwords from read-only memory describe by a buffer a constant (v#) through the constant cache (kcache).

```

m_offset = IMM ? OFFSET : SGPR[OFFSET]
m_base = { SGPR[SBASE * 2 + 1][15:0], SGPR[SBASE * 2] }
m_stride = SGPR[SBASE * 2 + 1][31:16]
m_num_records = SGPR[SBASE * 2 + 2]
m_size = (m_stride == 0) ? 1 : m_num_records
m_addr = (SGPR[SBASE * 2] + m_offset) & ~0x3
SGPR[SDST] = read_dword_from_kcache(m_base, m_offset, m_size)
SGPR[SDST + 1] = read_dword_from_kcache(m_base, m_offset + 4, m_size)
SGPR[SDST + 2] = read_dword_from_kcache(m_base, m_offset + 8, m_size)
. . .
SGPR[SDST + 7] = read_dword_from_kcache(m_base, m_offset + 28, m_size)
    
```

Microcode SMEM Opcode 11 (0xB)




Instruction **S_BUFFER_LOAD_DWORDX16**

Description Read 16 Dwords from read-only memory describe by a buffer a constant (v#) through the constant cache (kcache).

```

m_offset = IMM ? OFFSET : SGPR[OFFSET]
m_base = { SGPR[SBASE * 2 + 1][15:0], SGPR[SBASE * 2] }
m_stride = SGPR[SBASE * 2 + 1][31:16]
m_num_records = SGPR[SBASE * 2 + 2]
m_size = (m_stride == 0) ? 1 : m_num_records
m_addr = (SGPR[SBASE * 2] + m_offset) & ~0x3
SGPR[SDST] = read_dword_from_kcache(m_base, m_offset, m_size)
SGPR[SDST + 1] = read_dword_from_kcache(m_base, m_offset + 4, m_size)
SGPR[SDST + 2] = read_dword_from_kcache(m_base, m_offset + 8, m_size)
. . .
SGPR[SDST + 15] = read_dword_from_kcache(m_base, m_offset + 60, m_size)
    
```


Microcode SMEM Opcode 12 (0xC)



Instruction **S_BUFFER_STORE_DWORD**

Description Write one Dword to scalar data cache. See S_STORE_DWORD for details on the offset input.


Microcode SMEM Opcode 24 (0x18)



Instruction **S_BUFFER_STORE_DWORDX2**

Description Write two Dwords to scalar data cache. See S_STORE_DWORD for details on the offset input.


Microcode SMEM Opcode 25 (0x19)



Instruction **S_BUFFER_STORE_DWORDX4**

Description Write four Dwords to scalar data cache. See S_STORE_DWORD for details on the offset input.


Microcode SMEM Opcode 26 (0x1A)



Instruction **S_DCACHE_INV**

Description Invalidate entire L1 constant cache.


Microcode SMEM Opcode 32 (0x20)



Instruction **S_DCACHE_INV_VOL**

Description Invalidate all volatile lines in L1 constant cache.


Microcode SMEM Opcode 34 (0x22)



Instruction **S_DCACHE_WB**

Description Invalidate all volatile lines in L1 constant cache.


Microcode SMEM Opcode 33 (0x21)



Instruction **S_DCACHE_WB_VOL**

Description Write back dirty data in the scalar data cache volatile lines.

Microcode SMEM Opcode 35 (0x23)




Instruction **S_LOAD_DWORD**

Description Read one Dword from read-only constant memory through the constant cache (kcache).

$m_offset = IMM ? OFFSET : SGPR[OFFSET]$
 $m_addr = (SGPR[SBASE] + m_offset) \& \sim 0x3$
 $SGPR[SDST] = read_dword_from_kcache(m_addr)$

Microcode SMEM Opcode 0 (0x0)




Instruction **S_LOAD_DWORDX2**

Description Read two Dwords from read-only constant memory through the constant cache (kcache).

$m_offset = IMM ? OFFSET : SGPR[OFFSET]$
 $m_addr = (SGPR[SBASE * 2] + m_offset) \& \sim 0x3$
 $SGPR[SDST] = read_dword_from_kcache(m_addr)$
 $SGPR[SDST+1] = read_dword_from_kcache(m_addr+4)$

Microcode SMEM Opcode 1 (0x1)




Instruction **S_LOAD_DWORDX4**

Description Read four Dwords from read-only constant memory through the constant cache (kcache).

```

m_offset = IMM ? OFFSET : SGPR[OFFSET]
m_addr = (SGPR[SBASE * 2] + m_offset) & ~0x3
SGPR[SDST] = read_dword_from_kcache(m_addr)
SGPR[SDST+1] = read_dword_from_kcache(m_addr+4)
SGPR[SDST+2] = read_dword_from_kcache(m_addr+8)
SGPR[SDST+3] = read_dword_from_kcache(m_addr+12)
    
```

Microcode SMEM Opcode 2 (0x2)




Instruction **S_LOAD_DWORDX8**

Description Read eight Dwords from read-only constant memory through the constant cache (kcache).

```

m_offset = IMM ? OFFSET : SGPR[OFFSET]
m_addr = (SGPR[SBASE * 2] + m_offset) & ~0x3
SGPR[SDST] = read_dword_from_kcache(m_addr)
SGPR[SDST+1] = read_dword_from_kcache(m_addr+4)
SGPR[SDST+2] = read_dword_from_kcache(m_addr+8)
. . .
SGPR[SDST+7] = read_dword_from_kcache(m_addr+28)
    
```

Microcode SMEM Opcode 3 (0x3)




Instruction **S_LOAD_DWORDX16**

Description Read 16 Dwords from read-only constant memory through the constant cache (kcache).

```

m_offset = IMM ? OFFSET : SGPR[OFFSET]
m_addr = (SGPR[SBASE * 2] + m_offset) & ~0x3
SGPR[SDST] = read_dword_from_kcache(m_addr)
SGPR[SDST+1] = read_dword_from_kcache(m_addr+4)
SGPR[SDST+2] = read_dword_from_kcache(m_addr+8)
. . .
SGPR[SDST+15] = read_dword_from_kcache(m_addr+60)
    
```


Microcode SMEM Opcode 4 (0x4)



Instruction **S_MEMREALTIME**

Description Return current 64-bit RTC.


Microcode SMEM Opcode 37 (0x25)



Instruction **S_MEMTIME**

Description Return current 64-bit timestamp. This “time” is a free-running clock counter based on the shader core clock.


Microcode SMEM Opcode 36 (0x24)



Instruction **S_STORE_DWORD**

Description Write one Dword to scalar data cache. If the offset is specified as an SGPR, the SGPR contains an unsigned BYTE offset (the two LSBs are ignored). If the offset is specified as an immediate 20-bit constant, the constant is an unsigned byte offset.


Microcode SMEM Opcode 16 (0x10)



Instruction **S_STORE_DWORDX2**

Description Write two Dwords to scalar data cache. See S_STORE_DWORD for details on the offset input.


Microcode SMEM Opcode 17 (0x11)



Instruction **S_STORE_DWORDX4**

Description Write four Dwords to scalar data cache. See S_STORE_DWORD for details on the offset input.


Microcode SMEM Opcode 18 (0x12)



Instruction **S_STORE_DWORDX4**

Description Write four Dwords to scalar data cache. See S_STORE_DWORD for details on the offset input.

Microcode SMEM Opcode 18 (0x12)



12.7 VOP2 Instructions

Instruction **V_ADD_F16**

Description $D.f16 = S0.f16 + S1.f16$. Supports denormals, round mode, exception flags, saturation.

Microcode VOP2 Opcode 31 (0x1F)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Instruction **V_ADD_F32**

Description Floating-point add.
 $D.f = S0.f + S1.f$.

Microcode VOP2 Opcode 1 (0x1)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3a Opcode 257 (0x101)

	NEG	OMOD	SRC2	SRC1	SRC0	+4		
1	1	0	1	0	0			
	OP			CL MP	r	ABS	VDST	+0

Instruction **V_ADD_U16**

Description $D.u16 = S0.u16 + S1.u16$. Supports saturation (unsigned 16-bit integer domain).

Microcode VOP2 Opcode 38 (0x26)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Instruction **V_ADD_U32**

Description Integer add based on unsigned integer components. Note: This is the same operation for signed and unsigned integers.
 $D.u = S0.u + S1.u$; $VCC[threadId] = (S0.u + S1.u \geq 0x80000000ULL ? 1 : 0)$ is an UNSIGNED overflow or carry-out for V_ADDC_U32. In VOP3, the VCC destination can be an arbitrary SGPR-pair.

Microcode VOP2 Opcode 25 (0x19)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3b Opcode 281 (0x119)

1	1	0	1	0	0	OP	CL MP	SDST	VDST	+0
						SRC2	SRC1		SRC0	+4

Instruction **V_ADDC_U32**

Description Integer add based on unsigned integer components, with carry in. Produces a carry out in VCC or a scalar register.
 Output carry bit of unsigned integer ADD.
 $D.u = S0.u + S1.u + VCC$; $VCC=carry-out$ (VOP3:sgpr=carry-out, S2.u=carry-in).

Microcode VOP2 Opcode 28 (0x1C)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3b Opcode 284 (0x11C)

1	1	0	1	0	0	OP	CL MP	SDST	VDST	+0
						SRC2	SRC1		SRC0	+4

Instruction **V_AND_B32**

Description Logical bit-wise AND.
 $D.u = S0.u \& S1.u$. Input and output modifiers not supported.

Microcode VOP2 Opcode 19 (0x13)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3a Opcode 275 (0x113)

1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0	
NEG		OMOD		SRC2			SRC1			SRC0		+4

Instruction **V_ASHREV_I16**

Description 16-bit arithmetic shift right with arguments reversed.
 $D.i[15:0] = \text{signext}(S1.i[15:0]) \gg S0.i[3:0]$. The vacated bits are set to the sign bit of the input value. SQ translates this to an internal SP opcode.

Microcode VOP2 Opcode 44 (0x2C)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Instruction **V_ASHRREV_I32**

Description Arithmetic shift right with arguments reversed.
 $D.i = \text{signtext}(S1.i) \gg S0.i[4:0]$. The vacated bits are set to the sign bit of the input value.

Microcode VOP2 Opcode 17 (0x11)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3a Opcode 273 (0x111)

1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0
---	---	---	---	---	---	----	----------	---	-----	------	----

Instruction **V_CNDMASK_B32**

Description Boolean conditional based on bit mask from SGPRs or VCC.
 $D.u = VCC[i] ? S1.u : S0.u$ ($i = \text{threadID}$ in wave); VOP3: specify VCC as a scalar GPR in S2.

Microcode VOP2 Opcode 0 (0x0)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3a Opcode 256 (0x100)

1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0
---	---	---	---	---	---	----	----------	---	-----	------	----

Instruction **V_LDEXP_F16**

Description Load exponent.
 $D.f16 = S0.f16 * (2 ** S1.i16).$

Microcode VOP2 Opcode 51 (0x33)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Instruction **V_LSHLREV_B16**

Description 16-bit logical shift left with reversed arguments.
 $D.u[15:0] = S1.u[15:0] \ll S0.u[3:0].$ SQ translates this to an internal SP opcode.

Microcode VOP2 Opcode 42 (0x2A)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Instruction **V_LSHLREV_B32**

Description Logical shift left with reversed arguments. Note: Arguments are reversed because it is common for the shift value to be an immediate; this requires placing the immediate in the Src0 position.
 $D.u = S1.u \ll S0.u[4:0].$

Microcode VOP2 Opcode 18 (0x12)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3a Opcode 274 (0x112)

NEG	OMOD	SRC2	SRC1	SRC0	+4
1	1	0	1	0	0
OP			CLMP	r	ABS
			VDST		+0

Instruction **V_LSHRREV_B16**

Description 16-bit logical shift right with arguments reversed.
 $D.u[15:0] = S1.u[15:0] \gg S0.u[3:0]$. The vacated bits are set to zero. SQ translates this to an internal SP opcode.

Microcode VOP2 Opcode 43 (0x2B)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Instruction **V_LSHRREV_B32**

Description $D.u = S1.u \gg S0.u[4:0]$.

Microcode VOP2 Opcode 16 (0x10)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3a Opcode 278 (0x116)

NEG	OMOD	SRC2	SRC1	SRC0	+4		
1	1	0	1	0	0		
OP			CLMP	r	ABS	VDST	+0

Instruction **V_MAC_F16**

Description 16-bit floating point multiply -accumulate.
 $D.f16 = S0.f16 * S1.f16 + D.f16$. Supports round mode, exception flags, saturation.
 SQ translates this to V_MAD_F16.

Microcode VOP2 Opcode 35 (0x23)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Instruction **V_MAC_F32**

Description Floating-point multiply accumulate.
 $D.f = S0.f * S1.f + D.f.$

Microcode VOP2 Opcode 22 (0x16)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3a Opcode 272 (0x110)

1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0
						SRC2	SRC1			SRC0	+4

Instruction **V_MADAK_F16**

Description 16-bit floating-point multiply-add with constant add operand.
 $D.f16 = S0.f16 * S1.f16 + K.f16$; K is a 16-bit inline constant stored in the following literal Dword. This opcode cannot use the VOP3 encoding and cannot use input/output modifiers. Supports round mode, exception flags, saturation. SQ translates this to V_MAD_F16.


Microcode VOP2 Opcode 37 (0x25)

LITERAL										+4
0	OP	VDST	VSRC1							+0

Instruction **V_MADAK_F32**

Description $D.f = S0.f * S1.f + K$; K is a 32-bit literal constant.


Microcode VOP2 Opcode 24 (0x18)



Instruction **V_MADMK_F16**

Description 16-bit floating-point multiply-add with multiply operand immediate.
 $D.f16 = S0.f16 * K.f16 + S1.f16$; K is a 16-bit inline constant stored in the following literal Dword. This opcode cannot use the VOP3 encoding and cannot use input/output modifiers. Supports round mode, exception flags, saturation. SQ translates this to V_MAD_F16.


Microcode VOP2 Opcode 36 (0x24)



Instruction **V_MADMK_F32**

Description $D.f = S0.f * K + S1.f$; K is a 32-bit literal constant.

Microcode VOP2 Opcode 23 (0x17)



Instruction **V_MAX_F16**

Description D.f16 = max(S0.f16, S1.f16). IEEE compliant. Supports denormals, round mode, exception flags, saturation.

Microcode VOP2 Opcode 45 (0x2D)

LITERAL					+4
0	OP	VDST	VSR1	SRC0	+0

Instruction **V_MAX_F32**

Description

```

if (ieee_mode)
  if (S0.f==sNaN)      result = quiet(S0.f);
  else if (S1.f==sNaN) result = quiet(S1.f);
  else if (S0.f==NaN)  result = S1.f;
  else if (S1.f==NaN)  result = S0.f;
  else if (S0.f>S1.f)  result = S0.f;
  else                  result = S1.f;
else
  else if (S0.f==NaN)  result = S1.f;
  else if (S1.f==NaN)  result = S0.f;
  else if (S0.f>=S1.f) result = S0.f;
  else                  result = S1.f;
    
```

Microcode VOP2 Opcode 11 (0xB)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3a Opcode 267 (0x10B)

NEG	OMOD	SRC2	SRC1	SRC0	+4		
1	1	0	1	0	0		
OP			CL MP	r	ABS	VDST	+0

Instruction **v_MAX_I16**

Description $D.i[15:0] = \max(S0.i[15:0], S1.i[15:0]).$

Microcode VOP2 Opcode 48 (0x30)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Instruction **v_MAX_I32**

Description Integer maximum based on signed integer components.
 $D.i = \max(S0.i, S1.i).$

Microcode VOP2 Opcode 13 (0xD)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3a Opcode 269 (0x10D)

NEG	OMOD	SRC2	SRC1	SRC0	+4		
1	1	0	1	0	0		
OP			CL MP	r	ABS	VDST	+0

Instruction **v_MAX_U16**

Description $D.u[15:0] = \max(S0.u[15:0], S1.u[15:0]).$

Microcode VOP2 Opcode 47 (0x2F)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Instruction **V_MAX_U32**

Description Integer maximum based on unsigned integer components.

```
If (S0.u >= S1.u)
    D.u = S0.u;
Else
    D.u = S1.u;
```

Microcode VOP2 Opcode 15 (0xF)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3a Opcode 271 (0x10F)

1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0	
						SRC2	SRC1		SRC0			+4

Instruction **V_MBCNT_HI_U32_B32**

Description Masked bit count of the upper 32 threads (threads 32-63). For each thread, this instruction returns the number of active threads which come before it.
 ThreadMask = (1 << ThreadPosition) - 1;
 D.u = CountOneBits(S0.u & ThreadMask[63:32]) + S1.u.

Microcode VOP3a Opcode 653 (0x28D)

1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0	
						SRC2	SRC1		SRC0			+4

Instruction **V_MBCNT_LO_U32_B32**

Description Masked bit count set 32 low. ThreadPosition is the position of this thread in the wavefront (in 0..63).
 ThreadMask = (1 << ThreadPosition) - 1;
 D.u = CountOneBits(S0.u & ThreadMask[31:0]) + S1.u.

Microcode VOP3a Opcode 652 (0x28C)

NEG		OMOD		SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_MIN_F16**

Description D.f16 = min(S0.f16, S1.f16). IEEE compliant. Supports denormals, round mode, exception flags, saturation.

Microcode VOP2 Opcode 46 (0x2E)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Instruction **V_MIN_F32**

Description

```

if (ieee_mode)
  if (S0.f==NaN)    result = quiet(S0.f);
  else if (S1.f==NaN) result = quiet(S1.f);
  else if (S0.f==NaN) result = S1.f;
  else if (S1.f==NaN) result = S0.f;
  else if (S0.f<S1.f) result = S0.f;
  else              result = S1.f;
else
  else if (S0.f==NaN) result = S1.f;
  else if (S1.f==NaN) result = S0.f;
  else if (S0.f<S1.f) result = S0.f;
  else              result = S1.f;
    
```

Microcode VOP2 Opcode 10 (0xA)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3a Opcode 266 (0x10A)

1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0
---	---	---	---	---	---	----	----------	---	-----	------	----

Instruction **V_MIN_I16**

Description

$D.i[15:0] = \min(S0.i[15:0], S1.i[15:0]).$

Microcode VOP2 Opcode 50 (0x32)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Instruction **V_MIN_I32**

Description Integer minimum based on signed integer components.

```

If (S0.i < S1.i)
    D.i = S0.i;
Else
    D.i = S1.i;
    
```

Microcode VOP2 Opcode 12 (0xC)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3a Opcode 268 (0x10C)

NEG	OMOD	SRC2	SRC1	SRC0	+4		
1	1	0	1	0	0		
		OP	CL MP	r	ABS	VDST	+0

Instruction **V_MIN_U16**

Description $D.u[15:0] = \min(S0.u[15:0], S1.u[15:0]).$

Microcode VOP2 Opcode 49 (0x31)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Instruction **V_MIN_U32**

Description Integer minimum based on signed unsigned integer components.

If (S0.u < S1.u)
 D.u = S0.u;
 Else
 D.u = S1.u;

Microcode VOP2 Opcode 14 (0xE)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3a Opcode 270 (0x10E)

NEG	OMOD	SRC2	SRC1	SRC0	+4		
1	1	0	1	0	0		
		OP	CL MP	r	ABS	VDST	+0

Instruction **V_MUL_F16**

Description D.f16 = S0.f16 * S1.f16. Supports denormals, round mode, exception flags, saturation.

Microcode VOP2 Opcode 34 (0x22)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Instruction **V_MUL_F32**

Description Floating point multiply. Uses IEEE rules for 0*anything.
 $D.f = S0.f * S1.f.$

Microcode VOP2 Opcode 5 (0x5)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3a Opcode 262 (0x105)

NEG	OMOD	SRC2	SRC1	SRC0	+4
1	1	0	1	0	0
OP					
CL	MP	r	ABS	VDST	+0

Instruction **V_MUL_HI_I32_I24**

Description 24-bit signed integer multiply.
 S0 and S1 are treated as 24-bit signed integers. Bits [31:24] are ignored. The result represents the high-order 16 bits of the 48-bit multiply result, sign extended to 32 bits:
 $D.i = (S0.i[23:0] * S1.i[23:0]) \gg 32.$

Microcode VOP2 Opcode 7 (0x7)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3a Opcode 263 (0x107)

NEG	OMOD	SRC2	SRC1	SRC0	+4
1	1	0	1	0	0
OP					
CL	MP	r	ABS	VDST	+0

Instruction **V_MUL_HI_U32_U24**

Description 24-bit unsigned integer multiply.
 S0 and S1 are treated as 24-bit unsigned integers. Bits [31:24] are ignored. The result represents the high-order 16 bits of the 48-bit multiply result: {16'b0, mul_result[47:32]}.
 $D.i = (S0.u[23:0] * S1.u[23:0]) \gg 32.$

Microcode VOP2 Opcode 9 (0x9)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3a Opcode 265 (0x109)

1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0
						SRC2	SRC1			SRC0	+4

Instruction **V_MUL_I32_I24**

Description 24 bit signed integer multiply
 Src a and b treated as 24 bit signed integers. Bits [31:24] ignored. The result represents the low-order 32 bits of the 48 bit multiply result: mul_result[31:0].
 $D.i = S0.i[23:0] * S1.i[23:0].$

Microcode VOP2 Opcode 6 (0x6)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3a Opcode 262 (0x106)

1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0
						SRC2	SRC1			SRC0	+4

Instruction **V_MUL_LEGACY_F32**

Description Floating-point multiply.
 $D.f = S0.f * S1.f$ (DX9 rules, $0.0*x = 0.0$).

Microcode VOP2 Opcode 4 (0x4)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Microcode VOP3a Opcode 260 (0x104)

1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0	
NEG		OMOD		SRC2			SRC1			SRC0		+4

Instruction **V_MUL_LO_U16**

Description $D.u16 = S0.u16 * S1.u16$. Supports saturation (unsigned 16-bit integer domain).

Microcode VOP2 Opcode 41 (0x29)

0	OP	VDST	VSR1	SRC0	+0
---	----	------	------	------	----

Instruction **V_MUL_U32_U24**

Description 24-bit unsigned integer multiply.
 S0 and S1 are treated as 24-bit unsigned integers. Bits [31:24] are ignored. The result represents the low-order 32 bits of the 48-bit multiply result: mul_result[31:0].
 $D.u = S0.u[23:0] * S1.u[23:0]$.

Microcode VOP2 Opcode 8 (0x8)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3a Opcode 264 (0x108)

NEG	OMOD	SRC2	SRC1	SRC0	+4
1	1	0	1	0	0
OP					
CL	MP	r	ABS	VDST	+0

Instruction **V_OR_B32**

Description Logical bit-wise OR.
 $D.u = S0.u | S1.u$.

Microcode VOP2 Opcode 20 (0x14)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3a Opcode 276 (0x114)

NEG	OMOD	SRC2	SRC1	SRC0	+4
1	1	0	1	0	0
OP					
CL	MP	r	ABS	VDST	+0

Instruction **V_SUB_U32**

Description $D.u = S0.u - S1.u$; $VCC[threadId] = (S1.u > S0.u ? 1 : 0)$ is an unsigned overflow or carry-out for V_SUBB_U32. In VOP3, the VCC destination can be an arbitrary SGPR-pair.

Microcode VOP2 Opcode 26 (0x1A)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3b Opcode 282 (0x11A)

NEG	OMOD	SRC2	SRC1	SRC0	+4	
1	1	0	1	0	0	
OP			CL MP	SDST	VDST	+0

Instruction **V_SUBB_U32**

Description Integer subtract based on signed or unsigned integer components, with borrow in. Produces a borrow out in VCC or a scalar register.
 $D.u = S0.u - S1.u - VCC$; $VCC=carry-out$ (VOP3:sgpr=carry-out, S2.u=carry-in).

Microcode VOP2 Opcode 29 (0x1D)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3b Opcode 285 (0x11D)

NEG	OMOD	SRC2	SRC1	SRC0	+4	
1	1	0	1	0	0	
OP			CL MP	SDST	VDST	+0

Instruction **V_SUBBREV_U32**

Description $D.u = S1.u - S0.u - VCC$; VCC=carry-out (VOP3:sgpr=carry-out, S2.u=carry-in).

Microcode VOP2 Opcode 30 (0x1E)

0	OP	VDST	VSRC1	SRC0	+0
----------	----	------	-------	------	----

Microcode VOP3b Opcode 286 (0x11E)

NEG	OMOD	SRC2	SRC1	SRC0	+4					
1	1	0	1	0	0	OP	CL MP	SDST	VDST	+0

Instruction **V_SUBREV_F16**

Description $D.f16 = S1.f16 - S0.f16$. Supports denormals, round mode, exception flags, saturation.
SQ translates to V_ADD_F16.

Microcode VOP2 Opcode 33 (0x21)

0	OP	VDST	VSRC1	SRC0	+0
----------	----	------	-------	------	----

Instruction **V_SUBREV_F32**

Description $D.f = S1.f - S0.f.$

Microcode VOP2 Opcode 3 (0x3)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3a Opcode 259 (0x103)

NEG	OMOD	SRC2	SRC1	SRC0	+4
1	1	0	1	0	0
OP					
CL	MP	r	ABS	VDST	+0

Instruction **V_SUBREV_U16**

Description $D.u16 = S1.u16 - S0.u16.$ Supports saturation (unsigned 16-bit integer domain). SQ translates this to V_SUB_U16 with reversed operands.

Microcode VOP2 Opcode 40 (0x28)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Instruction **V_SUBREV_U32**

Description $D.u = S1.u - S0.u;$
 $VCC[threadId] = (S0.u > S1.u ? 1 : 0)$ is an unsigned overflow or carry-out for V_SUBB_U32 . In VOP3, the VCC destination can be an arbitrary SGPR-pair. SQ translates this to V_SUB_U32 with reversed operands.

Microcode VOP2 Opcode 27 (0x1B)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3b Opcode 283 (0x11B)

NEG	OMOD	SRC2	SRC1	SRC0	+4					
1	1	0	1	0	0	OP	CL MP	SDST	VDST	+0

Instruction **V_WRITELANE_B32**

Description Write value into one VGPR one one lane. Dst = VGPR-dest, Src0 = Source Data (SGPR, M0, exec, or constants), Src1 = Lane Select (SGPR or M0). Ignores exec mask.

Microcode VOP3a Opcode 650 (0x28A)

NEG	OMOD	SRC2	SRC1	SRC0	+4						
1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0

Instruction **V_XOR_B32**

Description Logical bit-wise XOR.
 $D.u = S0.u \wedge S1.u.$

Microcode VOP2 Opcode 21 (0x15)

0	OP	VDST	VSRC1	SRC0	+0
---	----	------	-------	------	----

Microcode VOP3a Opcode 277 (0x115)

NEG	OMOD	SRC2	SRC1	SRC0	+4
1	1	0	1	0	0
OP			CL MP	r	ABS
				VDST	+0

12.8 VOP1 Instructions

Instruction **V_BFREV_B32**

Description Bitfield reverse.
 $D.u[31:0] = S0.u[0:31]$.

Microcode VOP1 Opcode 44 (0x2C)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 364 (0x16C)

NEG	OMOD	SRC2				SRC1			SRC0			+4
1	1	0	1	0	0	OP	CLMP	r	ABS	VDST	+0	

Instruction **V_CEIL_F16**

Description Floating point ceiling function.
 $D.f16 = \text{trunc}(S0.f16)$; if $(S0.f16 > 0.0f \ \&\& \ S0.f16 \neq D.f16)$ then $D.f16 += 1.0f$.

Microcode VOP1 Opcode 69 (0x45)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_CEIL_F32**

Description Floating point ceiling function.
 D.f = ceil(S0.f). Implemented as: D.f = trunc(S0.f);
 if (S0 > 0.0 && S0 != D), D += 1.0.

Microcode VOP1 Opcode 29 (0x1D)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 349 (0x15D)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_CEIL_F64**

Description 64-bit floating-point ceiling.
 D.d = trunc(S0.d); if (S0.d > 0.0 && S0.d != D.d), D.d += 1.0.

Microcode VOP1 Opcode 24 (0x158)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 344 (0x15F)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0


Instruction **V_CLREXCP**

Description Clear wave's exception state in SIMD.

Microcode VOP1 Opcode 53 (0x35)



Microcode VOP3a Opcode 373 (0x175)



Instruction **V_COS_F16**

Description Cos function.
 Input must be normalized from radians by dividing by 2*PI.
 Valid input domain [-256, +256], which corresponds to an un-normalized input domain [-512*PI, +512*PI].
 Out-of-range input results in float 1.
 $D.f16 = \cos(S0.f16 * 2 * PI).$

Microcode VOP1 Opcode 74 (0x4A)



Instruction **v_cos_f32**

Description Cos function.
 Input must be normalized from radians by dividing by 2*PI.
 Valid input domain [-256, +256], which corresponds to an un-normalized input domain [-512*PI, +512*PI].
 Out-of-range input results in float 1.
 D.f = cos(S0.f).

Microcode VOP1 Opcode 42 (0x2A)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 362 (0x16A)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP	CLMP	r	ABS	VDST	+0		

Instruction **v_cvt_f16_f32**

Description Float32 to Float16.
 D.f16 = flt32_to_flt16(S0.f).

Microcode VOP1 Opcode 10 (0xA)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 330 (0x14A)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP	CLMP	r	ABS	VDST	+0		

Instruction **V_CVT_F16_I16**

Description D.f16 = int16_to_flt16(S.i16). Supports denormals, rounding, exception flags and saturation.

Microcode VOP1 Opcode 58 (0x3A)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_CVT_F16_U16**

Description D.f16 = uint16_to_flt16(S.u16). Supports denormals, rounding, exception flags and saturation.

Microcode VOP1 Opcode 57 (0x39)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_CVT_F32_F16**

Description DX11 Float16 to Float32.
D.f = flt16_to_flt32(S0.f16).

Microcode VOP1 Opcode 11 (0xB)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 331 (0x14B)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP		CLMP	r	ABS	VDST		+0

Instruction **V_CVT_F32_F64**

Description Convert Double Precision Float to Single Precision Float.
Overflows obey round mode rules. Infinity is exact.
D.f = (float)S0.d.

Microcode VOP1 Opcode 15 (0xF)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 335 (0x14F)

NEG		OMOD		SRC2				SRC1			SRC0		+4	
1	1	0	1	0	0	OP				CL MP	r	ABS	VDST	+0

Instruction **V_CVT_F32_I32**

Description The input is interpreted as a signed integer value and converted to a float.
D.f = (float)S0.i.

Microcode VOP1 Opcode 5 (0x5)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 325 (0x145)

NEG		OMOD		SRC2				SRC1			SRC0		+4	
1	1	0	1	0	0	OP				CL MP	r	ABS	VDST	+0

Instruction **V_CVT_F32_U32**

Description The input is interpreted as an unsigned integer value and converted to a float.
 $D.f = (\text{float})S0.u$.

Microcode VOP1 Opcode 6 (0x6)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 326 (0x146)

NEG	OMOD	SRC2					SRC1			SRC0	+4
1	1	0	1	0	0	OP	CLMP	r	ABS	VDST	+0

Instruction **V_CVT_F32_UBYTE0**

Description Byte 0 to float. Perform unsigned int to float conversion on byte 0 of S0.
 $D.f = \text{UINT2FLT}(S0.u[7:0])$.

Microcode VOP1 Opcode 17 (0x11)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 337 (0x151)

NEG	OMOD	SRC2					SRC1			SRC0	+4
1	1	0	1	0	0	OP	CLMP	r	ABS	VDST	+0

Instruction **V_CVT_F32_UBYTE1**

Description Byte 1 to float. Perform unsigned int to float conversion on byte 1 of S0.
 $D.f = \text{UINT2FLT}(S0.u[15:8]).$

Microcode VOP1 Opcode 18 (0x12)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 338 (0x152)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_CVT_F32_UBYTE2**

Description Byte 2 to float. Perform unsigned int to float conversion on byte 2 of S0.
 $D.f = \text{UINT2FLT}(S0.u[23:16]).$

Microcode VOP1 Opcode 19 (0x13)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 339 (0x153)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_CVT_F32_UBYTE3**

Description Byte 3 to float. Perform unsigned int to float conversion on byte 3 of S0.
 $D.f = \text{UINT2FLT}(S0.u[31:24]).$

Microcode VOP1 Opcode 20 (0x14)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 340 (0x154)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST	+0

Instruction **V_CVT_F64_F32**

Description Convert Single Precision Float to Double Precision Float.
 $D.d = (\text{double})S0.f.$

Microcode VOP1 Opcode 16 (0x10)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 336 (0x150)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST	+0

Instruction **V_CVT_F64_I32**

Description Convert Signed Integer to Double Precision Float.
 D.f = (float)S0.i.

Microcode VOP1 Opcode 4 (0x4)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 324 (0x144)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_CVT_F64_U32**

Description Convert Unsigned Integer to Double Precision Float.
 D.d = (double)S0.u.

Microcode VOP1 Opcode 22 (0x16)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 342 (0x156)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_CVT_FLR_I32_F32**

Description Float input is converted to a signed integer value using floor function. Float magnitudes too great to be represented by an integer float (unbiased exponent > 30) saturate to max_int or -max_int.
 $D.i = (\text{int})\text{floor}(S0.f)$.

Microcode VOP1 Opcode 13 (0xD)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 333 (0x14D)

NEG	OMOD	SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP	CLMP	r	ABS	VDST	+0

Instruction **V_CVT_I16_F16**

Description $D.i16 = \text{flt16_to_int16}(S.f16)$. Supports rounding, exception flags and saturation.

Microcode VOP1 Opcode 60 (0x3C)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_CVT_I32_F32**

Description Float input is converted to a signed integer using truncation.
 Float magnitudes too great to be represented by an integer float (unbiased exponent > 30) saturate to max_int or -max_int.
 Special case number handling:
 inf → max_int
 -inf → -max_int
 NaN & -NaN & 0 & -0 → 0
 D.i = (int)S0.f.

Microcode VOP1 Opcode 8 (0x8)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 328 (0x148)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP		CL MP	r	ABS	VDST		+0

Instruction **V_CVT_I32_F64**

Description Covert Double Precision Float to Signed Integer.
 Truncate (round-to-zero) only. Other round modes require a rne_f64, ceil_f64 or floor_f64 pre-op. Float magnitudes too great to be represented by an integer float (unbiased exponent > 30) saturate to max_int or -max_int.
 Special case number handling:
 inf → max_int
 -inf → -max_int
 NaN & -NaN & 0 & -0 → 0
 D.i = (int)S0.d.

Microcode VOP1 Opcode 3 (0x3)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 323 (0x143)


NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP		CL MP	r	ABS	VDST		+0

Instruction **V_CVT_OFF_F32_I4**


Description 4-bit signed int to 32-bit float. For interpolation in shader.

<u>S0</u>	<u>Result</u>
1000	-0.5f
1001	-0.4375f
1010	-0.375f
1011	-0.3125f
1100	-0.25f
1101	-0.1875f
1110	-0.125f
1111	-0.0625f
0000	0.0f
0001	0.0625f
0010	0.125f
0011	0.1875f
0100	0.25f
0101	0.3125f
0110	0.375f
0111	0.4375f

Microcode VOP1 Opcode 14 (0xE)



Microcode VOP3a Opcode 334 (0x14E)



Instruction **V_CVT_RPI_I32_F32**


Description Float input is converted to a signed integer value using round to positive infinity tiebreaker for 0.5. Float magnitudes too great to be represented by an integer float (unbiased exponent > 30) saturate to max_int or -max_int.

$$D.i = (\text{int})\text{floor}(S0.f + 0.5).$$

Microcode VOP1 Opcode 12 (0xC)




Microcode VOP3a Opcode 332 (0x14C)



Instruction **V_CVT_U16_F16**

Description $D.u16 = \text{flt16_to_uint16}(S.f16)$. Supports rounding, exception flags and saturation.

Microcode VOP1 Opcode 59 (0x3B)




Instruction **V_CVT_U32_F32**

Description Input is converted to an unsigned integer value using truncation. Positive float magnitudes too great to be represented by an unsigned integer float (unbiased exponent > 31) saturate to max_uint.
 Special number handling:
 -inf & NaN & 0 & -0 → 0
 Inf → max_uint
 D.u = (unsigned)S0.f.

Microcode VOP1 Opcode 7 (0x7)



Microcode VOP3a Opcode 327 (0x147)




Instruction **V_CVT_U32_F64**

Description Covert Double Precision Float to Unsigned Integer
 Truncate (round-to-zero) only. Other round modes require a rne_f64, ceil_f64 or floor_f64 pre-op. Positive float magnitudes too great to be represented by an unsigned integer float (unbiased exponent > 31) saturate to max_uint.
 Special number handling:
 -inf & NaN & 0 & -0 → 0
 Inf → max_uint
 D.u = (uint)S0.d.

Microcode VOP1 Opcode 21 (0x15)



Microcode VOP3a Opcode 341 (0x155)




Instruction **V_EXP_F32**

Description Base2 exponent function.


```

If (Arg1 == 0.0f) {
    Result = 1.0f;
}
Else {
    Result = Approximate2ToX(Arg1);
}
    
```

Microcode VOP1 Opcode 32 (0x20)



Microcode VOP3a Opcode 352 (0x160)




Instruction **V_EXP_F16**

Description Base2 exponent function.

```

if(S0.f16 == 0.0f) D.f16 = 1.0f; else D.f16 = Approximate2ToX(S0.f16).
    
```

Microcode VOP1 Opcode 65 (0x41)



Instruction **V_EXP_LEGACY_F32**

Description Return 2^{argument} floating-point value, using the same precision as GCN Generation 1.
 D.f = pow(2.0, S0.f). Same as GCN Generation 1.

Microcode VOP1 Opcode 75 (0x4B)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 395 (0x18B)

NEG	OMOD	SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP	CLMP	r	ABS	VDST	+0

Instruction **V_FFBH_I32**

Description Find first bit signed high.
 Find first bit set in a positive integer from MSB, or find first bit clear in a negative integer from MSB
 D.u = position of first bit different from sign bit in S0 from MSB;
 D=0xFFFFFFFF if S0==0 or 0xFFFFFFFF.

Microcode VOP1 Opcode 47 (0x2F)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 367 (0x16F)

NEG	OMOD	SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP	CLMP	r	ABS	VDST	+0

Instruction **V_FFBH_U32**

Description Find first bit high.
 D.u = position of first 1 in S0 from MSB; D=0xFFFFFFFF if S0==0.

Microcode VOP1 Opcode 45 (0x2D)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 365 (0x16D)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_FFBL_B32**

Description Find first bit low.
 D.u = position of first 1 in S0 from LSB; D=0xFFFFFFFF if S0==0.

Microcode VOP1 Opcode 46 (0x2E)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 366 (0x16E)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_FLOOR_F16**

Description Floating-point floor function.
 $D.f16 = \text{trunc}(S0.f16); \text{ if}(S0.f16 < 0.0f \ \&\& \ S0.f16 \neq D.f16) \text{ then } D.f16 -= 1.0f.$

Microcode VOP1 Opcode 68 (0x44)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_FLOOR_F32**

Description Floating-point floor function.
 $D.f = \text{trunc}(S0); \text{ if}((S0 < 0.0) \ \&\& \ (S0 \neq D)) \ D += -1.0.$

Microcode VOP1 Opcode 31 (0x1F)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 351 (0x15F)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	0	0	OP	CLMP	r	ABS	VDST	+0

Instruction **V_FLOOR_F64**

Description 64-bit floating-point floor.
 $D.d = \text{trunc}(S0.d); \text{ if } (S0.d < 0.0 \ \&\& \ S0.d \neq D.d), \ D.d += -1.0.$

Microcode VOP1 Opcode 26 (0x1A)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 346 (0x15A)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP	CLMP	r	ABS	VDST	+0		

Instruction **V_FRACT_F16**

Description Floating point 'fractional' part of S0.f.
 $D.f16 = S0.f16 + -\text{floor}(S0.f16).$

Microcode VOP1 Opcode 72 (0x48)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_FRACT_F32**

Description Floating point 'fractional' part of S0.f.
 $D.f = S0.f - \text{floor}(S0.f)$.

Microcode VOP1 Opcode 27 (0x1B)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 347 (0x15B)

NEG		OMOD		SRC2				SRC1			SRC0		+4	
1	1	0	1	0	0	OP				CLMP	r	ABS	VDST	+0

Instruction **V_FRACT_F64**

Description Double-precision fractional part of Arg1.
 Double result written to two consecutive GPRs; the instruction Dest specifies the lesser of the two.
 $D.d = \text{FRAC64}(S0.d)$.
 Return fractional part of input as double [0.0 - 1.0).

Microcode VOP1 Opcode 50 (0x32)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 370 (0x172)

NEG		OMOD		SRC2				SRC1			SRC0		+4	
1	1	0	1	0	0	OP				CLMP	r	ABS	VDST	+0

Instruction **V_FREXP_EXP_I16_F16**

Description if(S0.f16 == +-INF || S0.f16 == NAN) D.i16 = 0; else D.i16 = 2_s_complement(exponent(S0.f16) - 15 + 1). C math library frexp function. Returns exponent of half precision float input, such that the original single float = significand * (2 ** exponent).

Microcode VOP1 Opcode 67 (0x43)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_FREXP_EXP_I32_F32**

Description C math library frexp function. Returns the exponent of a single precision float input, such that:
 original single float = significand * 2^{exponent}.
 D.f = 2's complement (exponent(S0.f) - 127 + 1).

Microcode VOP1 Opcode 51 (0x33)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 371 (0x173)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_FREXP_EXP_I32_F64**

Description C++ FREXP math function.
Returns exponent of double precision float input, such that:
original double float = significand * 2^{exponent}.

D.i = 2's complement (exponent(S0.d) - 1023 +1).

Microcode VOP1 Opcode 48 (0x30)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 368 (0x170)

NEG	OMOD	SRC2				SRC1			SRC0			+4
1	1	0	1	0	0	OP	CLMP	r	ABS	VDST	+0	

Instruction **V_FREXP_MANT_F16**

Description if(S0.f16 == +-INF || S0.f16 == NAN) D.f16 = S0.f16; else D.f16 = mantissa(S0.f16). Result range is (-1.0,-0.5][0.5,1.0). C math library frexp function. Returns binary significand of half precision float input, such that the original single float = significand * (2 ** exponent).

Microcode VOP1 Opcode 66 (0x42)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_FREXP_MANT_F32**

Description C math library frexp function. Returns binary significand of single precision float input, such that:
 original single float = significand * 2^{exponent} .
 D.f =Mantissa(S0.f).
 D.f range(-1.0,-0.5] or [0.5,1.0).

Microcode VOP1 Opcode 52 (0x34)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 372 (0x174)

NEG	OMOD	SRC2	SRC1	SRC0	+4						
1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0

Instruction **V_FREXP_MANT_F64**

Description C++ FREXP math function.
 Returns binary significand of double precision float input, such that
 original double float = significand * 2^{exponent} .
 D.d =Mantissa(S0.d).
 D.d range(-1.0,-0.5] or [0.5,1.0).

Microcode VOP1 Opcode 49 (0x31)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 369 (0x171)

NEG	OMOD	SRC2	SRC1	SRC0	+4						
1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0

Instruction **V_LOG_F16**

Description Base2 log function.
 if(S0.f16 == 1.0f) D.f16 = 0.0f; else D.f16 = ApproximateLog2(S0.f16).

Microcode VOP1 Opcode 64 (0x40)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_LOG_F32**

Description Base2 log function.
 D.f = log2(S0.f).

Microcode VOP1 Opcode 33 (0x21)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 353 (0x161)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST	+0

Instruction **V_LOG_LEGACY_F32**

Description Return the algorithm of a 32-bit floating point value, using the same precision as GCN Generation 1.
 $D.f = \log_2(S0.f)$. Base 2 logarithm. Same as GCN Generation 1.

Microcode VOP1 Opcode 76 (0x4C)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 396 (0x18C)

NEG		OMOD		SRC2				SRC1			SRC0		+4	
1	1	0	1	0	0	OP				CL MP	r	ABS	VDST	+0

Instruction **V_MOV_B32**

Description Single operand move instruction. Allows denorms in and out, regardless of denorm mode, in both single and double precision designs.
 $D.u = S0.u$.

Microcode VOP1 Opcode 1 (0x1)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 321 (0x141)

NEG		OMOD		SRC2				SRC1			SRC0		+4	
1	1	0	1	0	0	OP				CL MP	r	ABS	VDST	+0

Instruction **V_MOVRELD_B32**

Description $VGPR[D.u + M0.u] = VGPR[S0.u]$.

Microcode VOP1 Opcode 54 (0x36)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 374 (0x176)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST	+0

Instruction **V_MOVRELS_B32**

Description $VGPR[D.u] = VGPR[S0.u + M0.u]$.

Microcode VOP1 Opcode 55 (0x37)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 375 (0x177)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST	+0

Instruction **V_MOVRELSD_B32**

Description $VGPR[D.u + M0.u] = VGPR[S0.u + M0.u]$.

Microcode VOP1 Opcode 56 (0x38)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 376 (0x178)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_NOP**

Description Do nothing.

Microcode VOP1 Opcode 0 (0x0)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 320 (0x140)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_NOT_B32**

Description Logical bit-wise NOT.
 D.u = ~S0.u.

Microcode VOP1 Opcode 43 (0x2B)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 363 (0x16B)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP		CL MP	r	ABS	VDST		+0

Instruction **V_RCP_F16**

Description if (S0.f16 == 1.0f), D.f16 = 1.0f; else D.f16 = ApproximateRecip(S0.f16).

Microcode VOP1 Opcode 61 (0x3D)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_RCP_F32**

Description Reciprocal, < 1 ulp error.
 This reciprocal approximation converges to < 0.5 ulp error with one Newton-Raphson performed with two fused multiple adds (FMAs).
 D.f = 1.0 / S0.f.

Microcode VOP1 Opcode 34 (0x22)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 354 (0x162)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_RCP_F64**

Description Double reciprocal.
 Inputs from two consecutive GPRs, the instruction source specifies lower of the two. Double result written to two consecutive GPRs; the instruction Dest specifies the lesser of the two.
 D.d = 1.0 / (S0.d).

Microcode VOP1 Opcode 37 (0x25)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 357 (0x165)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_RCP_IFLAG_F32**

Description $D.f = 1.0 / src0.f$. Reciprocal intended for integer division, can raise integer DIV_BY_ZERO exception but cannot raise floating-point exceptions.

Microcode VOP1 Opcode 35 (0x23)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_READFIRSTLANE_B32**

Description Copy one VGPR value to one SGPR. Dst = SGPR-dest, Src0 = Source Data (VGPR# or M0(lds-direct)), Lane# = FindFirst1fromLSB(exec) (lane = 0 if exec is zero). Executes regardless of exec mask value.

Microcode VOP1 Opcode 2 (0x2)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 322 (0x142)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP		CLMP	r	ABS	VDST		+0

Instruction **V_RNDNE_F16**

Description Floating-point Round-to-Nearest-Even Integer.
 $D.f16 = \text{FLOOR}(S0.f16 + 0.5f)$; if (floor(src0.f16) is even && fract(src0.f16) == 0.5f) then $D.f16 -= 1.0f$. Round-to-nearest-even semantics.

Microcode VOP1 Opcode 71 (0x47)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_RNDNE_F32**

Description Floating-point Round-to-Nearest-Even Integer.
 $D.f = \text{round_nearest_even}(S0.f)$.

Microcode VOP1 Opcode 30 (0x1E)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 350 (0x15e)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0		

Instruction **V_RNDNE_F64**

Description 64-bit floating-point round-to-nearest-even.
 $D.d = \text{round_nearest_even}(S0.d)$.

Microcode VOP1 Opcode 25 (0x19)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 345 (0x159)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0		

Instruction **V_RSQ_F16**

Description $\text{if}(S0.f16 == 1.0f) \text{ D.f16} = 1.0f; \text{ else } \text{D.f16} = \text{ApproximateRecipSqrt}(S0.f16).$


Microcode VOP1 Opcode 63 (0x3F)




Instruction **V_RSQ_F32**

Description Reciprocal square roots.
 $\text{D.f} = 1.0 / \text{sqrt}(S0.f).$

Microcode VOP1 Opcode 36 (0x24)



Microcode VOP3a Opcode 356 (0x164)



Instruction **V_RSQ_F64**

Description Double reciprocal square root.
 Inputs from two consecutive GPRs; the instruction source specifies the lower of the two registers. The double result is written to two consecutive GPRs; the instruction Dest specifies the lesser of the two.
 $D.f = 1.0 / \text{sqrt}(S0.f)$.

Microcode VOP1 Opcode 38 (0x26)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 358 (0x166)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST	+0

Instruction **V_SIN_F16**

Description Sin function.
 Input must be normalized from radians by dividing by 2*PI.
 Valid input domain [-256, +256], which corresponds to an un-normalized input domain [-512*PI, +512*PI].
 Out of range input results in float 0.
 $D.f16 = \sin(S0.f16 * 2 * \text{PI})$.

Microcode VOP1 Opcode 73 (0x49)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----


Instruction **V_SIN_F32**

Description Sin function.
 Input must be normalized from radians by dividing by 2π .
 Valid input domain $[-256, +256]$, which corresponds to an un-normalized input domain $[-512\pi, +512\pi]$.
 Out of range input results in float 0.
 $D.f = \sin(S0.f)$.

Microcode VOP1 Opcode 41 (0x29)




Microcode VOP3a Opcode 361 (0x169)



Instruction **V_SQRT_F16**

Description $\text{if}(S0.f16 == 1.0f) D.f16 = 1.0f; \text{ else } D.f16 = \text{ApproximateSqrt}(S0.f16)$.

Microcode VOP1 Opcode 62 (0x2E)



Instruction **V_SQRT_F32**

Description Square root. Useful for normal compression.
 $D.f = \text{sqrt}(S0.f)$.

Microcode VOP1 Opcode 39 (0x27)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 359 (0x167)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_SQRT_F64**

Description $D.d = \text{sqrt}(S0.d)$.

Microcode VOP1 Opcode 40 (0x28)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 360 (0x168)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_TRUNC_F16**

Description Floating point 'integer' part of S0.f.
 D.f16 = trunc(S0.f16). Round-to-zero semantics.

Microcode VOP1 Opcode 70 (0x46)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Instruction **V_TRUNC_F32**

Description Floating point 'integer' part of S0.f.
 D.f = trunc(S0.f), return integer part of S0.

Microcode VOP1 Opcode 28 (0x1C)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 348 (0x15c)

NEG		OMOD		SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+0

Instruction **V_TRUNC_F64**

Description Truncate a 64-bit floating-point value, and return the resulting integer value.
 $D.d = \text{trunc}(S0.d)$, return integer part of $S0.d$.

Microcode VOP1 Opcode 23 (0x17)

0	1	1	1	1	1	1	1	VDST	OP	SRC0	+0
---	---	---	---	---	---	---	---	------	----	------	----

Microcode VOP3a Opcode 343 (0x157)

NEG		OMOD		SRC2				SRC1			SRC0		+4	
1	1	0	1	0	0	OP				CL MP	r	ABS	VDST	+0

12.9 VOPC Instructions


The bitfield map for VOPC is:



where:

- SRC0 = First operand for instruction.
- VSRC1 = Second operand for instruction.
- OP = Instructions.

All VOPC instructions are also part of VOP3a microcode format, for which the bitfield is:



where:

- VDST = Destination for instruction in the VGPR.
- ABS = Floating-point absolute value.
- CLMP = Clamp output.
- OP = Instructions.
- SRC0 = First operand for instruction.
- SRC1 = Second operand for instruction.
- SRC2 = Third operand for instruction. Unused in VOPC instructions.
- OMOD = Output modifier for instruction. Unused in VOPC instructions.
- NEG = Floating-point negation.

The first eight VOPC instructions have {OP16} embedded in them. This refers to each of the compare operations listed below.

<u>Compare Operation</u>	<u>Opcode Offset</u>	<u>Description</u>
F	0	D.u = 0
LT	1	D.u = (S0 < S1)
EQ	2	D.u = (S0 == S1)
LE	3	D.u = (S0 <= S1)
GT	4	D.u = (S0 > S1)
LG	5	D.u = (S0 <> S1)
GE	6	D.u = (S0 >= S1)
O	7	D.u = (!isNaN(S0) && !isNaN(S1))
U	8	D.u = (!isNaN(S0) !isNaN(S1))
NGE	9	D.u = !(S0 >= S1)

NLG	10	D.u = !(S0 <> S1)
NGT	11	D.u = !(S0 > S1)
NLE	12	D.u = !(S0 <= S1)
NEQ	13	D.u = !(S0 == S1)
NLT	14	D.u = !(S0 < S1)
TRU	15	D.u = 1

Table 12.1 VOPC Instructions with 16 Compare Operations

Instruction	Description	Hex Range
V_CMP_{OP16}_F32	Signal on signalling NaN (sNaN) input only.	0x00 to 0x0F
V_CMPX_{OP16}_F32	Signal on sNaN input only. Also write EXEC.	0x10 to 0x1F
V_CMP_{OP16}_F64	Signal on sNaN input only.	0x20 to 0x2F
V_CMPX_{OP16}_F64	Signal on sNaN input only. Also write EXEC.	0x30 to 0x3F
V_CMPS_{OP16}_F32	Signal on any NaN.	0x40 to 0x4F
V_CMPSX_{OP16}_F32	Signal on any NaN. Also write EXEC.	0x50 to 0x5F
V_CMPS_{OP16}_F64	Signal on any NaN.	0x60 to 0x6F
V_CMPSX_{OP16}_F64	Signal on any NaN. Also write EXEC.	0x70 to 0x7F

The second eight VOPC instructions have {OP8} embedded in them. This refers to each of the compare operations listed below.

Compare Operation	Opcode Offset	Description
F	0	D.u = 0
LT	1	D.u = (S0 < S1)
EQ	2	D.u = (S0 == S1)
LE	3	D.u = (S0 <= S1)
GT	4	D.u = (S0 > S1)
LG	5	D.u = (S0 <> S1)
GE	6	D.u = (S0 >= S1)
TRU	7	D.u = 1

Table 12.2 VOPC Instructions with Eight Compare Operations

Instruction	Description	Hex Range
V_CMP_{OP8}_I32	On 32-bit integers.	0x80 to 0x87
V_CMPX_{OP8}_I32	Also write EXEC.	0x90 to 0x97
V_CMP_{OP8}_I64	On 64-bit integers.	0xA0 to 0xA7
V_CMPX_{OP8}_I64	Also write EXEC.	0xB0 to 0xB7
V_CMP_{OP8}_U32	On unsigned 32-bit integers.	0xC0 to 0xC7
V_CMPX_{OP8}_U32	Also write EXEC.	0xD0 to 0xD7
V_CMP_{OP8}_U64	On unsigned 64-bit integers.	0xE0 to 0xE7
V_CMPX_{OP8}_U64	Also write EXEC.	0xF0 to 0xF7

The final instructions for VOPC are four CLASS instructions.

Table 12.3 VOPC CLASS Instructions

Instruction	Description	Hex
V_CMP_CLASS_F32	D = IEEE numeric class function specified in S1.u, performed on S0.f.	0x10
V_CMPX_CLASS_F32	D = IEEE numeric class function specified in S1.u, performed on S0.f. Also write EXEC.	0x11
V_CMP_CLASS_F64	D = IEEE numeric class function specified in S1.u, performed on S0.d. Result is single bit Boolean for each thread, aggregated across wavefront and returned to SQ. Result is true if Arg1 is a member of any of the classes indicated by the mask (Arg2). mask[0] - signalingNaN mask[1] - quietNaN mask[2] - negativeInfinity mask[3] - negativeNormal mask[4] - negativeSubnormal mask[5] - negativeZero mask[6] - positiveZero mask[7] - positiveSubnormal mask[8] - positiveNormal mask[9] - positiveInfinity There is no vector result written to a gpr, and no vector feedback path for this opcode. This opcode does not raise exceptions under any circumstances.	0x12
V_CMPX_CLASS_F64	D = IEEE numeric class function specified in S1.u, performed on S0.d. Also write EXEC. Result is single bit Boolean for each thread, aggregated across wavefront and returned to SQ. Result is true if Arg1 is a member of any of the classes indicated by the mask (Arg2). mask[0] - signalingNaN mask[1] - quietNaN mask[2] - negativeInfinity mask[3] - negativeNormal mask[4] - negativeSubnormal mask[5] - negativeZero mask[6] - positiveZero mask[7] - positiveSubnormal mask[8] - positiveNormal mask[9] - positiveInfinity There is no vector result written to a gpr, and no vector feedback path for this opcode. This opcode does not raise exceptions under any circumstances.	0x13
V_COMP_CLASS_F16	D = IEEE numeric class function specified in S1.u, performed on S0.f.	0x14
V_CMPX_CLASS_F16	D = IEEE numeric class function specified in S1.u, performed on S0.f. Also write EXEC.	0x15

12.10 VOP3 3 in, 1 out Instructions (VOP3a)

Add Floating-Point, 64-Bit

Instruction **V_ADD_F64**

Description Double-precision floating-point add.
 Floating-point 64-bit add. Adds two double-precision numbers in the YX or WZ elements of the source operands, src0 and src1, and outputs a double-precision value to the same elements of the destination operand. No carry or borrow beyond the 64-bit values is performed. The operation occupies two slots in an instruction group. Double result written to two consecutive GPR registers, instruction dest specifies lower of the two.
 $D.d = S0.d + S1.d$

Table 12.4 Result of V_ADD_F64 Instruction

src0	src1								
	-inf	-F ¹	-denorm	-0	+0	+denorm	+F ¹	+inf	NaN ²
-inf	-inf	-inf	-inf	-inf	-inf	-inf	-inf	NaN64	src1 (NaN64)
-F ¹	-inf	-F	src0	src0	src0	src0	+-F or +0	+inf	src1 (NaN64)
-denorm	-inf	src1	-0	-0	+0	+0	src1	+inf	src1 (NaN64)
-0	-inf	src1	-0	-0	+0	+0	src1	+inf	src1 (NaN64)
+0	-inf	src1	+0	+0	+0	+0	src1	+inf	src1 (NaN64)
+denorm	-inf	src1	+0	+0	+0	+0	src1	+inf	src1 (NaN64)
+F ¹	-inf	+-F or +0	src0	src0	src0	src0	+F	+inf	src1 (NaN64)
+inf	NaN64	+inf	+inf	+inf	+inf	+inf	+inf	+inf	src1 (NaN64)
NaN	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)

1. F is a finite floating-point value.
2. NaN64 = 0xFFF8000000000000. An NaN64 is a propagated NaN value from the input listed.

Add Floating-Point, 64-Bit (Cont.)

These properties hold true for this instruction:
 $(A + B) == (B + A)$
 $(A - B) == (A + -B)$
 $A + -A = +zero$

Microcode VOP3a Opcode 640 (0x280)

NEG			OMOD			SRC2			SRC1			SRC0			
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST		+4	
													+0		

Instruction **V_ALIGNBIT_B32**

Description Bit align. Arbitrarily align 32 bits within 64 into a GPR.
 $D.u = (\{S0,S1\} \gg S2.u[4:0]) \& 0xFFFFFFFF.$

Microcode VOP3a Opcode 462 (0x1CE)

NEG			OMOD			SRC2			SRC1			SRC0			
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST		+4	
													+0		

Instruction **V_ALIGNBYTE_B32**

Description Byte align.
 $dst = (\{src0, src1\} \gg (8 * src2[1:0])) \& 0xFFFFFFFF;$
 $D.u = (\{S0,S1\} \gg (8*S2.u[4:0])) \& 0xFFFFFFFF.$

Microcode VOP3a Opcode 463 (0x1CF)

NEG			OMOD			SRC2			SRC1			SRC0			
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST		+4	
													+0		

Instruction **V_ASHRREV_I64**

Description D.u64 = signext(S1.u64) >> S0.u[5:0]. The vacated bits are set to the sign bit of the input value.

Microcode VOP3a Opcode 657 (0x291)

NEG			OMOD			SRC2						SRC1			SRC0			+4
1	1	0	1	0	0	OP						CL MP	r	ABS	VDST			+0

Instruction **V_BCNT_U32_B32**

Description Bit count.
D.u = CountOneBits(S0.u) + S1.u.

Microcode VOP3a Opcode 651 (0x28B)

NEG			OMOD			SRC2						SRC1			SRC0			+4
1	1	0	1	0	0	OP						CL MP	r	ABS	VDST			+0

Instruction **V_BFE_I32**

Description DX11 signed bitfield extract. src0 = input data, src1 = offset, and src2 = width. The bit position offset is extracted through offset + width from the input data. All bits remaining after dst are stuffed with replications of the sign bit.

```

If (src2[4:0] == 0) {
    dst = 0;
}
Else if (src2[4:0] + src1[4:0] < 32) {
    dst = (src0 << (32-src1[4:0] - src2[4:0])) >>> (32 - src2[4:0])
}
Else {
    dst = src0 >>> src1[4:0]
}

```

D.i = (S0.i>>S1.u[4:0]) & ((1<<S2.u[4:0])-1); bitfield extract, S0=data, S1=field_offset, S2=field_width.

Microcode VOP3a Opcode 457 (0x1C9)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_BFE_U32**

Description DX11 unsigned bitfield extract. Src0 = input data, src1 = offset, and src2 = width. Bit position offset is extracted through offset + width from input data.

```

If (src2[4:0] == 0) {
    dst = 0;
}
Else if (src2[4:0] + src1[4:0] < 32) {
    dst = (src0 << (32-src1[4:0] - src2[4:0])) >> (32 - src2[4:0])
}
Else {
    dst = src0 >> src1[4:0]
}

```

D.u = (S0.u>>S1.u[4:0]) & ((1<<S2.u[4:0])-1); bitfield extract, S0=data, S1=field_offset, S2=field_width.

Microcode VOP3a Opcode 456 (0x1C8)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_BFI_B32**

Description Bitfield insert used after BFM to implement DX11 bitfield insert.
 src0 = bitfield mask (from BFM)
 src 1 & src2 = input data
 This replaces bits in src2 with bits in src1 according to the bitfield mask.
 $D.u = (S0.u \& S1.u) \mid (\sim S0.u \& S2.u).$

Microcode VOP3a Opcode 458 (0x1CA)

NEG		OMOD		SRC2				SRC1			SRC0		+4	
1	1	0	1	0	0	OP				CL MP	r	ABS	VDST	+0

Instruction **V_BFM_B32**

Description Bitfield mask. Used before BFI to implement DX11 bitfield insert.
 $D.u = ((1 \ll S0.u[4:0]) - 1) \ll S1.u[4:0];$ S0=bitfield_width,
 S1=bitfield_offset.

Microcode VOP3a Opcode 659 (0x293)

NEG		OMOD		SRC2				SRC1			SRC0		+4	
1	1	0	1	0	0	OP				CL MP	r	ABS	VDST	+0

Instruction **V_CUBEID_F32**

Description Cubemap Face ID determination. Result is a floating point face ID.

```

S0.f = x
S1.f = y
S2.f = z
If (Abs(S2.f) >= Abs(S0.f) &&
    Abs(S2.f) >= Abs(S1.f))
    If (S2.f < 0) D.f = 5.0
    Else D.f = 4.0
Else if (Abs(S1.f) >= Abs(S0.f))
    If (S1.f < 0) D.f = 3.0
    Else D.f = 2.0
Else
    If (S0.f < 0) D.f = 1.0
    Else D.f = 0.0
    
```

Microcode VOP3a Opcode 452 (0x1C4)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_CUBE_MAJ_F32**

Description Cubemap Major Axis determination. Result is 2.0 * Major Axis.

```

S0.f = x
S1.f = y
S2.f = z
If (Abs(S2.f) >= Abs(S0.f) &&
    Abs(S2.f) >= Abs(S1.f))
    D.f = 2.0*S2.f
Else if (Abs(S1.f) >= Abs(S0.f))
    D.f = 2.0 * S1.f
Else
    D.f = 2.0 * S0.f
    
```

Microcode VOP3a Opcode 455 (0x1C7)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_CUBESC_F32**

Description Cubemap S coordination determination.

```

S0.f = x
S1.f = y
S2.f = z
If (Abs(S2.f) >= Abs(S0.f) &&
    Abs(S2.f) >= Abs(S1.f))
    If (S2.f < 0) D.f = -S0.f
    Else D.f = S0.f
Else if (Abs(S1.f) >= Abs(S0.f))
    D.f = S0.f
Else
    If (S0.f < 0) D.f = S2.f
    Else D.f = -S2.f
    
```

Microcode VOP3a Opcode 453 (0x1C5)

NEG			OMOD			SRC2						SRC1			SRC0			+4
1	1	0	1	0	0	OP						CL MP	r	ABS	VDST			+0

Instruction **V_CUBETC_F32**

Description Cubemap T coordinate determination.

```

S0.f = x
S1.f = y
S2.f = z
If (Abs(S2.f) >= Abs(S0.f) &&
    Abs(S2.f) >= Abs(S1.f))
    D.f = -S1.f
Else if (Abs(S1.f) >= Abs(S0.f))
    If (S1.f < 0) D.f = -S2.f
    Else D.f = S2.f
Else
    D.f = -S1.f
    
```

Microcode VOP3a Opcode 454 (0x1C6)

NEG			OMOD			SRC2						SRC1			SRC0			+4
1	1	0	1	0	0	OP						CL MP	r	ABS	VDST			+0

Instruction **V_CVT_PK_I16_I32**

Description DX signed 32-bit integer to signed 16-bit integer.
 Overflow clamped to 0x7FFF. Underflow clamped to 0x8000.
 $D = \{(i32@i16)S1.i, (i32@i16)S0.i\}.$

Microcode VOP3a Opcode 664 (0x298)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_CVT_PK_U8_F32**

Description Float to 8 bit unsigned integer conversion
 Replacement for 8xx/9xx FLT_TO_UINT4 opcode.
 Float to 8 bit uint conversion placed into any byte of result, accumulated with S2.f. Four applications of this opcode can accumulate 4 8-bit integers packed into a single dword.
 $D.f = ((flt_to_uint(S0.f) \& 0xff) \ll (8*S1.f[1:0])) \mid (S2.f \& \sim(0xff \ll (8*S1.f[1:0])));$

Intended use, ops in any order:
 op - cvt_pk_u8_f32 r0 foo2, 2, r0
 op - cvt_pk_u8_f32 r0 foo1, 1, r0
 op - cvt_pk_u8_f32 r0 foo3, 3, r0
 op - cvt_pk_u8_f32 r0 foo0, 0, r0

r0 result is 4 bytes packed into a Dword:
 {foo3, foo2, foo1, foo0}

Microcode VOP3a Opcode 477 (0x1DD)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_CVT_PK_U16_U32**

Description DX11 unsigned 32-bit integer to unsigned 16-bit integer.
 Overflow clamped to 0xFFFF.
 $D = \{(u32@u16)S1.u, (u32@u16)S0.u\}.$

Microcode VOP3a Opcode 663 (0x297)

NEG			OMOD			SRC2						SRC1						SRC0						+4
1	1	0	1	0	0	OP						CL	MP	r	ABS			VDST						+0

Instruction **V_CVT_PKACCU8_U8_F32**

Description $f32@u8(s0.f),$ pack into byte(s1.u), of dst.

Microcode VOP3a Opcode 496 (0x1F0)

NEG			OMOD			SRC2						SRC1						SRC0						+4
1	1	0	1	0	0	OP						CL	MP	r	ABS			VDST						+0

Instruction **V_CVT_PKNORM_I16_F32**

Description DX Float32 to SNORM16, a signed, normalized 16-bit value.
 $D = \{(snorm)S1.f, (snorm)S0.f\}.$

Microcode VOP3a Opcode 660 (0x294)

NEG			OMOD			SRC2						SRC1						SRC0						+4
1	1	0	1	0	0	OP						CL	MP	r	ABS			VDST						+0

Instruction **V_CVT_PKNORM_U16_F32**

Description DX Float32 to UNORM16, an unsigned, normalized 16-bit value.
 $D = \{(\text{unorm})S1.f, (\text{unorm})S0.f\}$.

Microcode VOP3a Opcode 661 (0x295)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_CVT_PKRTZ_F16_F32**

Description Convert two float 32 numbers into a single register holding two packed 16-bit floats.
 $D = \{\text{flt32_to_flt16}(S1.f), \text{flt32_to_flt16}(S0.f)\}$, with round-toward-zero.

Microcode VOP3a Opcode 662 (0x296)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_DIV_FIXUP_F16**

Description Given a numerator, denominator, and quotient from a divide, this opcode detects and applies special case numerics, modifies the quotient if necessary. This opcode also generates invalid, denorm, and divide by zero exceptions caused by the division.

```

sign_out = sign(S1.f16)^sign(S2.f16);
if (S2.f16 == NAN)
    D.f16 = Quiet(S2.f16);
else if (S1.f16 == NAN)
    D.f16 = Quiet(S1.f16);
else if (S1.f16 == S2.f16 == 0)
    # 0/0
    D.f16 = pele_nan(0xfe00);
else if (abs(S1.f16) == abs(S2.f16) == +-INF)
    # inf/inf
    D.f16 = pele_nan(0xfe00);
else if (S1.f16 == 0 || abs(S2.f16) == +-INF)
    # x/0, or inf/y
    D.f16 = sign_out ? -INF : INF;
else if (abs(S1.f16) == +-INF || S2.f16 == 0)
    # x/inf, 0/y
    D.f16 = sign_out ? -0 : 0;
else if ((exp(S2.f16) - exp(S1.f16)) < -150)
    D.f16 = sign_out ? -underflow : underflow;
else if (exp(S1.f16) == 255)
    D.f16 = sign_out ? -overflow : overflow;
else
    D.f16 = sign_out ? -abs(S0.f16) : abs(S0.f16).
Half precision division fixup.
S0 = Quotient, S1 = Denominator, S3 = Numerator.
    
```

Microcode VOP3a Opcode 495 (0x1EF)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_DIV_FIXUP_F32**

Description Single precision division fixup.
 Given a numerator, denominator, and quotient from a divide, this opcode detects and applies special-case numerics, touching up the quotient if necessary. This opcode also generates all exceptions caused by the division. The generation of the inexact exception requires a fused multiple add (FMA), making this opcode a variant of FMA.
 S0.f = Quotient
 S1.f = Denominator
 S2.f = Numerator
 If (S1.f==Nan && S2.f!=SNan)
 D.f = Quiet(S1.f);
 Else if (S2.f==Nan)
 D.f = Quiet(S2.f);
 Else if (S1.f==S2.f==0)
 # 0/0
 D.f = pele_nan(0xffc00000);
 Else if (abs(S1.f)==abs(S2.f)==infinity)
 # inf/inf
 D.f = pele_nan(0xffc00000);
 Else if (S1.f==0)
 # x/0
 D.f = (sign(S1.f)^sign(S0.f) ? -inf : inf);
 Else if (abs(S1.f)==inf)
 # x/inf
 D.f = (sign(S1.f)^sign(S0.f) ? -0 : 0);
 Else if (S0.f==Nan)
 # division error correction nan due to N*1/D overflow (result of divide is overflow)
 D.f = (sign(S1.f)^sign(S0.f) ? -inf : inf);
 Else
 D.f = S0.f;

Microcode VOP3a Opcode 478 (0x1DE)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_DIV_FIXUP_F64**

Description Double precision division fixup.
 Given a numerator, denominator, and quotient from a divide, this opcode detects and applies special case numerics, touching up the quotient if necessary. This opcode also generates all exceptions caused by the division. The generation of the inexact exception requires a fused multiply add (FMA), making this opcode a variant of FMA.
 D.d = Special case divide fixup and flags(s0.d = Quotient, s1.d = Denominator, s2.d = Numerator).

Microcode VOP3a Opcode 479 (0x1DF)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_DIV_FMAS_F32**

Description D.f = Special case divide FMA with scale and flags(s0.f = Quotient, s1.f = Denominator, s2.f = Numerator).

Microcode VOP3a Opcode 482 (0x1E2)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_DIV_FMAS_F64**

Description D.d = Special case divide FMA with scale and flags (s0.d = Quotient, s1.d = Denominator, s2.d = Numerator).

Microcode VOP3a Opcode 483 (0x1E3)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_FMA_F16**

Description Fused half precision multiply add.
 $D.f16 = S0.f16 * S1.f16 + S2.f16.$

Microcode VOP3a Opcode 494 (0x1EE)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_FMA_F32**

Description Fused single-precision multiply-add. Only for double-precision parts.
 $D.f = S0.f * S1.f + S2.f.$

Microcode VOP3a Opcode 459 (0x1CB)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_FMA_F64**

Description Double-precision floating-point fused multiply add (FMA).
 Adds the src2 to the product of the src0 and src1. A single round is performed on the sum - the product of src0 and src1 is not truncated or rounded.
 The instruction specifies which one of two data elements in a four-element vector is operated on (the two dwords of a double precision floating point number), and the result can be stored in the wz or yx elements of the destination GPR.
 $D.d = S0.d * S1.d + S2.d.$

Microcode VOP3a Opcode 460 (0x1CC)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_LDEXP_F32**

Description C math library ldexp function.
 Result = S0.f * (2 ^ S1.i)
 S0 = float 32
 S1 = signed integer

Microcode VOP3a Opcode 648 (0x288)

NEG			OMOD			SRC2						SRC1				SRC0				+4
1	1	0	1	0	0	OP						CL MP	r	ABS		VDST				+0

Instruction **V_LDEXP_F64**

Description Double-precision LDEXP from the C math library.
 This instruction gets a 52-bit mantissa from the double-precision floating-point value in src1.YX and a 32-bit integer exponent in src0.X, and multiplies the mantissa by 2^{exponent}. The double-precision floating-point result is stored in dst.YX.

```
dst = src1 * 2^src0

mant = mantissa(src1)
exp = exponent(src1)
sign = sign(src1)

if (exp==0x7FF) //src1 is inf or a NaN
{
    dst = src1;
}
else if (exp==0x0) //src1 is zero or a denorm
{
    dst = (sign) ? 0x8000000000000000 : 0x0;
}
else //src1 is a float
{
    exp+= src0;
    if (exp>=0x7FF) //overflow
    {
        dst = {sign,inf};
    }
    if (src0<=0) //underflow
    {
        dst = {sign,0};
    }

    mant |= (exp<<52);
    mant |= (sign<<63);

    dst = mant;
}
```

Table 12.5 Result of LDEXP_F64 Instruction

src1	src0				
	-/+inf	-/+denorm	-/+0	-/+F ¹	NaN
-/+I ²	-/+inf	-/+0	-/+0	src1 * (2 ^{src0})	src0
Not -/+I	-/+inf	-/+0	-/+0	invalid result	src0

1. F is a finite floating-point value.
2. I is a valid 32-bit integer value.

Microcode VOP3a Opcode 644 (0x284)

NEG		OMOD		SRC2			SRC1			SRC0			
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST	+4
												+0	

Instruction **V_LERP_U8**

Description Unsigned eight-bit pixel average on packed unsigned bytes (linear interpolation). S2 acts as a round mode; if set, 0.5 rounds up; otherwise, 0.5 truncates.

$$D.u = ((S0.u[31:24] + S1.u[31:24] + S2.u[24]) >> 1) << 24 + ((S0.u[23:16] + S1.u[23:16] + S2.u[16]) >> 1) << 16 + ((S0.u[15:8] + S1.u[15:8] + S2.u[8]) >> 1) << 8 + ((S0.u[7:0] + S1.u[7:0] + S2.u[0]) >> 1).$$

$$dst = ((src0[31:24] + src1[31:24] + src2[24]) >> 1) << 24 + ((src0[23:16] + src1[23:16] + src2[16]) >> 1) << 16 + ((src0[15:8] + src1[15:8] + src2[8]) >> 1) << 8 + ((src0[7:0] + src1[7:0] + src2[0]) >> 1)$$

Microcode VOP3a Opcode 461 (0x1CD)

NEG		OMOD		SRC2			SRC1			SRC0			
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST	+4
												+0	

Instruction **V_LSHLREV_B64**

Description $D.u64 = S1.u64 \ll S0.u[5:0]$.

Microcode VOP3a Opcode 655 (0x28F)

NEG		OMOD		SRC2				SRC1			SRC0			+4
1	1	0	1	0	0	OP				CL MP	r	ABS	VDST	+0

Instruction **V_LSHLRREV_B64**

Description $D.u64 = S1.u64 \gg S0.u[5:0]$. The vacated bits are set to zero.

Microcode VOP3a Opcode 656 (0x290)

NEG		OMOD		SRC2				SRC1			SRC0			+4
1	1	0	1	0	0	OP				CL MP	r	ABS	VDST	+0

Instruction **V_MAD_F16**

Description Floating point multiply-add (MAD). Gives same result as ADD after MUL_IEEE. Uses IEEE rules for 0*anything.
 $D.f16 = S0.f16 * S1.f16 + S2.f16$. Supports round mode, exception flags, saturation.

Microcode VOP3a Opcode 490 (0x1EA)

NEG		OMOD		SRC2				SRC1			SRC0			+4
1	1	0	1	0	0	OP				CL MP	r	ABS	VDST	+0

Instruction **V_MAD_F32**

Description Floating point multiply-add (MAD). Gives same result as ADD after MUL_IEEE. Uses IEEE rules for 0*anything.
 $D.f = S0.f * S1.f + S2.f.$

Microcode VOP3a Opcode 449 (0x1C1)

NEG		OMOD		SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST	+0

Instruction **V_MAD_I16**

Description Signed integer muladd.
 S0 and S1 are treated as 16-bit signed integers. S2 is treated as a 16-bit signed or unsigned integer. Bits [31:16] are ignored. The result represents the low-order sign extended 16 bits of the multiply add result.
 $D.i16 = S0.i16 * S1.i16 + S2.i16.$ Supports saturation (signed 16-bit integer domain).

Microcode VOP3a Opcode 492 (0x1EC)

NEG		OMOD		SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST	+0

Instruction **V_MAD_I32_I24**

Description 24-bit signed integer muladd.
 S0 and S1 are treated as 24-bit signed integers. S2 is treated as a 32-bit signed or unsigned integer. Bits [31:24] are ignored. The result represents the low-order sign extended 32 bits of the multiply add result.
 Result = Arg1.i[23:0] * Arg2.i[23:0] + Arg3.i[31:0] (low order bits).

Microcode VOP3a Opcode 450 (0x1C2)

NEG			OMOD			SRC2						SRC1				SRC0				+4
1	1	0	1	0	0	OP						CL MP	r	ABS	VDST				+0	

Instruction **V_MAD_I64_I32**

Description Multiply add using the product of two 32-bit signed integers, then added to a 64-bit integer.
 {vcc_out,D.i64} = S0.i32 * S1.i32 + S2.i64.

Microcode VOP3a Opcode 489 (0x1E9)

NEG			OMOD			SRC2						SRC1				SRC0				+4
1	1	0	1	0	0	OP						CL MP	r	ABS	VDST				+0	

Instruction **V_MAD_LEGACY_F32**

Description Floating-point multiply-add (MAD). Gives same result as ADD after MUL.
 D.f = S0.f * S1.f + S2.f (DX9 rules, 0.0*x = 0.0).

Microcode VOP3a Opcode 448 (0x1C0)

NEG			OMOD			SRC2						SRC1				SRC0				+4
1	1	0	1	0	0	OP						CL MP	r	ABS	VDST				+0	

Instruction **V_MAD_U16**

Description Unsigned integer muladd.
 $D.u16 = S0.u16 * S1.u16 + S2.u16$. Supports saturation (unsigned 16-bit integer domain).

Microcode VOP3a Opcode 491 (0x1EB)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_MAD_U32_U24**

Description 24 bit unsigned integer muladd
 Src a and b treated as 24 bit unsigned integers. Src c treated as 32 bit signed or unsigned integer. Bits [31:24] ignored. The result represents the low-order 32 bits of the multiply add result.
 $D.u = S0.u[23:0] * S1.u[23:0] + S2.u[31:0]$.

Microcode VOP3a Opcode 451 (0x1C3)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_MAD_U64_U32**

Description Multiply add using the product of two 32-bit unsigned integers, then added to a 64-bit integer.
 $\{vcc_out, D.u64\} = S0.u32 * S1.u32 + S2.u64$.

Microcode VOP3a Opcode 488 (0x1E8)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_MAX_F64**

Description The instruction specifies which one of two data elements in a four-element vector is operated on (the two Dwords of a double precision floating point number), and the result can be stored in the wz or yx elements of the destination GPR.

D.d = max(S0.d, S1.d).

```
if (src0 > src1)
    dst = src0;
else
    dst = src1;
```

max(-0,+0)=max(+0,-0)=+0

The rules for NaN handling for F_MAX_F64 are:

```
if (ieee_mode)
    if (Arg1==sNaN result = quiet(Arg1);
    else if (Arg2==sNaN result = quiet(Arg2);
    else if (Arg1==NaN) result = Arg2;
    else if (Arg2==NaN) result = Arg1;
    else if (Arg1>Arg2) result = Arg1;
    else result = Arg2;
else
    else if (Arg1==NaN) result = Arg2;
    else if (Arg2==NaN) result = Arg1;
    else if (Arg1>=Arg2) result = Arg1;
    else
        result = Arg2;
```

Microcode VOP3a Opcode 643 (0x283)

NEG			OMOD			SRC2			SRC1			SRC0				
1	1	0	1	0	0	OP			CL	r	ABS	VDST			+4	
												MP				+0

Instruction **V_MAX3_F32**

Description Maximum of three numbers. DX10 NaN handling and and flag creation.

D.f = max(S0.f, S1.f, S2.f).

Microcode VOP3a Opcode 467 (0x1D3)

NEG			OMOD			SRC2			SRC1			SRC0				
1	1	0	1	0	0	OP			CL	r	ABS	VDST			+4	
												MP				+0

Instruction **V_MAX3_I32**

Description Maximum of three numbers.
 $D.i = \max(S0.i, S1.i, S2.i)$.

Microcode VOP3a Opcode 468 (0x1D4)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_MAX3_U32**

Description Maximum of three numbers.
 $D.u = \max(S0.u, S1.u, S2.u)$.

Microcode VOP3a Opcode 469 (0x1D5)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_MBCNT_LO_U32_B32**

Description Masked bit count set 32 low. ThreadPosition is the position of this thread in the wavefront (in 0..63).
 $\text{ThreadMask} = (1 \ll \text{ThreadPosition}) - 1$;
 $D.u = \text{CountOneBits}(S0.u \& \text{ThreadMask}[31:0]) + S1.u$.

Microcode VOP3a Opcode 652 (0x28C)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_MED3_F32**

Description Median of three numbers. DX10 NaN handling and flag creation.

```

If (isNaN(S0.f) || isNaN(S1.f) || isNaN(S2.f))
    D.f = MIN3(S0.f, S1.f, S2.f)
Else if (MAX3(S0.f,S1.f,S2.f) == S0.f)
    D.f = MAX(S1.f, S2.f)
Else if (MAX3(S0.f,S1.f,S2.f) == S1.f)
    D.f = MAX(S0.f, S2.f)
Else
    D.f = MAX(S0.f, S1.f)
    
```

Microcode VOP3a Opcode 470 (0x1D6)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_MED3_I32**

Description Median of three numbers.

```

IF (MAX3(S0.i,S1.i,S2.i) == S0.i)
    D.i = MAX(S1.i, S2.i)
Else if (MAX3(S0.i,S1.i,S2.i) == S1.i)
    D.i = MAX(S0.i, S2.i)
Else
    D.i = MAX(S0.i, S1.i)
    
```

Microcode VOP3a Opcode 471 (0x1D7)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_MED3_U32**

Description Median of three numbers.

```

If (MAX3(S0.i,S1.i,S2.i) == S0.i)
    D.i = MAX(S1.i, S2.i)
Else if (MAX3(S0.i,S1.i,S2.i) == S1.i)
    D.i = MAX(S0.i, S2.i)
Else
    D.i = MAX(S0.i, S1.i)
    
```

Microcode VOP3a Opcode 472 (0x1D8)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_MIN_F64**

Description Double precision floating point minimum.
 The instruction specifies which one of two data elements in a four-element vector is operated on (the two dwords of a double precision floating point number), and the result can be stored in the wz or yx elements of the destination GPR.
 DX10 implies slightly different handling of Nan's. See the SP Numeric spec for details.
 Double result written to two consecutive GPRs; the instruction Dest specifies the lesser of the two.

```

if (src0 < src1)
    dst = src0;
else
    dst = src1;

min(-0,+0)=min(+0,-0)=-0

D.d = min(S0.d, S1.d).
    
```

Microcode VOP3a Opcode 642 (0x282)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_MIN3_F32**

Description Minimum of three numbers. DX10 NaN handling and flag creation.
 $D.f = \min(S0.f, S1.f, S2.f)$.

Microcode VOP3a Opcode 464 (0x1D0)

NEG			OMOD			SRC2						SRC1						SRC0						+4
1	1	0	1	0	0	OP						CL	MP	r	ABS			VDST			+0			

Instruction **V_MIN3_I32**

Description Minimum of three numbers.
 $D.i = \min(S0.i, S1.i, S2.i)$.

Microcode VOP3a Opcode 465 (0x1D1)

NEG			OMOD			SRC2						SRC1						SRC0						+4
1	1	0	1	0	0	OP						CL	MP	r	ABS			VDST			+0			

Instruction **V_MIN3_U32**

Description Minimum of three numbers.
 $D.u = \min(S0.u, S1.u, S2.u)$.

Microcode VOP3a Opcode 466 (0x1D2)

NEG			OMOD			SRC2						SRC1						SRC0						+4
1	1	0	1	0	0	OP						CL	MP	r	ABS			VDST			+0			

Instruction **V_MQSD_PK_U16_U8**

Description D.u = Masked Quad-Byte SAD with accum_lo/hi(S0.u[63:0], S1.u[31:0], S2.u[63:0]).

Microcode VOP3a Opcode 486 (0x1E6)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_MQSD_U32_U8**

Description Masked quad sum-of-absolute-difference.
 D.u128 = Masked Quad-Byte SAD with 32-bit accum_lo/hi(S0.u[63:0], S1.u[31:0], S2.u[127:0]).

Microcode VOP3a Opcode 487 (0x1E7)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_MSAD_U8**

Description D.u = Masked Byte SAD with accum_lo(S0.u, S1.u, S2.u).

Microcode VOP3a Opcode 484 (0x1E4)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_MUL_F64**

Description Floating-point 64-bit multiply. Multiplies a double-precision value in `src0.YX` by a double-precision value in `src1.YX`, and places the lower 64 bits of the result in `dst.YX`. Inputs are from two consecutive GPRs, with the instruction specifying the lesser of the two; the double result is written to two consecutive GPRs.
 $dst = src0 * src1;$
 $D.d = S0.d * S1.d$

Table 12.6 Result of MUL_64 Instruction

src0	src1										
	-inf	-F ¹	-1.0	-denorm	-0	+0	+denorm	+1.0	+F ¹	+inf	NaN ²
-inf	+inf	+inf	+inf	NaN64	NaN64	NaN64	NaN64	-inf	-inf	-inf	src1 (NaN64)
-F	+inf	+F	-src0	+0	+0	-0	-0	src0	-F	-inf	src1 (NaN64)
-1.0	+inf	-src1	+1.0	+0	+0	-0	-0	-1.0	-src1	-inf	src1 (NaN64)
-denorm	NaN64	+0	+0	+0	+0	-0	-0	-0	-0	NaN64	src1 (NaN64)
-0	NaN64	+0	+0	+0	+0	-0	-0	-0	-0	NaN64	src1 (NaN64)
+0	NaN64	-0	-0	-0	-0	+0	+0	+0	+0	NaN64	src1 (NaN64)
+denorm	NaN64	-0	-0	-0	-0	+0	+0	+0	+0	NaN64	src1 (NaN64)
+1.0	-inf	src1	-1.0	-0	-0	+0	+0	+1.0	src1	+inf	src1 (NaN64)
+F	-inf	-F	-src0	-0	-0	+0	+0	src0	+F	+inf	src1 (NaN64)
+inf	-inf	-inf	-inf	NaN64	NaN64	NaN64	NaN64	+inf	+inf	+inf	src1 (NaN64)
NaN	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)	src0 (NaN64)

1. F is a finite floating-point value.
2. NaN64 = 0xFFF8000000000000. An NaN64 is a propagated NaN value from the input listed.

$$(A * B) == (B * A)$$

Coissue The `V_MUL_F64` instruction is a four-slot instruction. Therefore, a single `V_MUL_F64` instruction can be issued in slots 0, 1, 2, and 3. Slot 4 can contain any other valid instruction.

Microcode VOP3a Opcode 641 (0x281)

NEG			OMOD			SRC2				SRC1				SRC0				+4
1	1	0	1	0	0	OP				CL	r	ABS	VDST				+0	
										MP								

Instruction **V_MUL_HI_I32**

Description Signed integer multiplication. The result represents the high-order 32 bits of the multiply result.
 $D.i = (S0.i * S1.i) \gg 32.$

Microcode VOP3a Opcode 647 (0x287)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_MUL_HI_U32**

Description Unsigned integer multiplication. The result represents the high-order 32 bits of the multiply result.
 $D.u = (S0.u * S1.u) \gg 32.$

Microcode VOP3a Opcode 646 (0x286)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_MUL_LO_U32**

Description Unsigned integer multiplication. The result represents the low-order 32 bits of the multiply result.
 $D.u = S0.u * S1.u.$

Microcode VOP3a Opcode 645 (0x285)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CLMP	r	ABS	VDST			+0

Instruction **V_PERM_B32**

Description **Byte permute.**
D.u[31:24] = permute({S0.u, S1.u}, S2.u[31:24]);
D.u[23:16] = permute({S0.u, S1.u}, S2.u[23:16]);
D.u[15:8] = permute({S0.u, S1.u}, S2.u[15:8]);
D.u[7:0] = permute({S0.u, S1.u}, S2.u[7:0]);
byte permute(byte in[8], byte sel) {
 if(sel>=13) then return 0xff;
 elseif(sel==12) then return 0x00;
 elseif(sel==11) then return in[7][7] * 0xff;
 elseif(sel==10) then return in[5][7] * 0xff;
 elseif(sel==9) then return in[3][7] * 0xff;
 elseif(sel==8) then return in[1][7] * 0xff;
 else return in[sel];
}

Microcode VOP3a Opcode 493 (0x1ED)

NEG		OMOD		SRC2			SRC1			SRC0			
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+4
												+0	

Instruction **V_QSAD_PK_U16_U8**

Description D.u = Quad-Byte SAD with accum_lo/hiu(S0.u[63:0], S1.u[31:0], S2.u[63:0]).

Microcode VOP3a Opcode 485 (0x1E5)

NEG		OMOD		SRC2			SRC1			SRC0			
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST	+4
												+0	

Instruction **V_READLANE_B32**

Description Copy one VGPR value to one SGPR. Dst = SGPR-dest, Src0 = Source Data (VGPR# or M0(lds-direct)), Src1 = Lane Select (SGPR or M0). Ignores exec mask. A lane corresponds to one thread in a wavefront.
The VOP3a version does not apply input or output modifiers.

Microcode VOP3a Opcode 649 (0x289)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_READLANE_B32**

Description Copy one VGPR value to one SGPR. Dst = SGPR-dest, Src0 = Source Data (VGPR# or M0(lds-direct)), Src1 = Lane Select (SGPR or M0). Ignores exec mask. A lane corresponds to one thread in a wavefront.

Microcode VOP3a Opcode 649 (0x289)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_SAD_HI_U8**

Description Sum of absolute differences with accumulation.
 Perform 4x1 SAD with S0.u and S1.u, and accumulate result into MSBs of S2.u. Overflow is lost.

$$\text{ABS_DIFF}(A,B) = (A>B) ? (A-B) : (B-A)$$

$$\begin{aligned} D.u = & (\text{ABS_DIFF}(S0.u[31:24],S1.u[31:24]) + \text{ABS_DIFF}(S0.u[23:16],S1.u[23:16]) \\ & + \text{ABS_DIFF}(S0.u[15:8],S1.u[15:8]) + \text{ABS_DIFF}(S0.u[7:0],S1.u[7:0])) \ll 16 + \\ & S2.u \end{aligned}$$

Microcode VOP3a Opcode 474 (0x1DA)

NEG			OMOD			SRC2						SRC1				SRC0				+4
1	1	0	1	0	0	OP						CL	MP	r	ABS	VDST				+0

Instruction **V_SAD_U8**

Description Sum of absolute differences with accumulation.
 Perform 4x1 SAD with S0.u and S1.u, and accumulate result into lsbs of S2.u. Overflow into S2.u upper bits is allowed.

$$\text{ABS_DIFF}(A,B) = (A>B) ? (A-B) : (B-A)$$

$$\begin{aligned} D.u = & \text{ABS_DIFF}(S0.u[31:24],S1.u[31:24]) + \text{ABS_DIFF}(S0.u[23:16],S1.u[23:16]) + \\ & \text{ABS_DIFF}(S0.u[15:8],S1.u[15:8]) + \text{ABS_DIFF}(S0.u[7:0],S1.u[7:0]) + S2.u \end{aligned}$$

Microcode VOP3a Opcode 473 (0x1D9)

NEG			OMOD			SRC2						SRC1				SRC0				+4
1	1	0	1	0	0	OP						CL	MP	r	ABS	VDST				+0

Instruction **V_SAD_U16**

Description Sum of absolute differences with accumulation.
Perform 2x1 SAD with S0.u and S1.u, and accumulate result with S2.u.

$$\text{ABS_DIFF}(A,B) = (A>B) ? (A-B) : (B-A)$$

$$D.u = \text{ABS_DIFF}(S0.u[31:16],S1.u[31:16]) + \text{ABS_DIFF}(S0.u[15:0],S1.u[15:0]) + S2.u$$

Microcode VOP3a Opcode 475 (0x1DB)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_SAD_U32**

Description Sum of absolute differences with accumulation.
Perform a single-element SAD with S0.u and S1.u, and accumulate result into MSB's of S2.u. Overflow is lost.

$$\text{ABS_DIFF}(A,B) = (A>B) ? (A-B) : (B-A)$$

$$D.u = \text{ABS_DIFF}(S0.u,S1.u) + S2.u$$

Microcode VOP3a Opcode 476 (0x1DC)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST			+0

Instruction **V_TRIG_PREOP_F64**

Description D.d = Look Up 2/PI (S0.d) with segment select S1.u[4:0].

Microcode VOP3a Opcode 658 (0x292)

NEG		OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	r	ABS	VDST		+0

12.11 VOP3 Instructions (3 in, 2 out), (VOP3b)

Instruction **V_DIV_SCALE_F32**

Description D.f = Special case divide preop and flags(s0.f = Quotient, s1.f = Denominator, s2.f = Numerator); s0 must equal s1 or s2.

Microcode VOP3b Opcode 480 (0x1E0)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	SDST			VDST	+0	

Instruction **V_DIV_SCALE_F64**

Description D.d = Special case divide preop and flags(s0.d = Quotient, s1.d = Denominator, s2.d = Numerator) s0 must equal s1 or s2.

Microcode VOP3b Opcode 481 (0x1E1)

NEG			OMOD			SRC2			SRC1			SRC0			+4
1	1	0	1	0	0	OP			CL MP	SDST			VDST	+0	

12.12 VINTRP Instructions

Instruction **V_INTERP_MOV_F32**

Description Vertex Parameter Interpolation using parameters stored in LDS and barycentric coordinates in VGPRs. D.f = Parameter value (p0, p10, p20) for this primitive.
M0 must contain: { 1'b0, new_prim_mask[15:1], lds_param_offset[15:0] }.
The ATTR field indicates which attribute (0-32) to interpolate.
The ATTRCHAN field indicates which channel: 0=x, 1=y, 2=z and 3=w.

Microcode VINTRP Opcode 2 (0x2)

1	1	0	1	0	1	VDST	OP	ATTR	ATTR CHAN	VSRC (I, J)	+0
---	---	---	---	---	---	------	----	------	--------------	-------------	----

Microcode VOP3a Opcode 626 (0x272)

NEG	OMOD	SRC2	SRC1	SRC0	+4						
1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0

Instruction **V_INTERP_P1_F32**

Description Vertex Parameter Interpolation using parameters stored in LDS and barycentric coordinates in VGPRs. First step for interpolation: D.f = P10.f * S0.f + P0, where P0 and P10 are parameters for this primitive from LDS.
M0 must contain: { 1'b0, new_prim_mask[15:1], lds_param_offset[15:0] }.
The ATTR field indicates which attribute (0-32) to interpolate.
The ATTRCHAN field indicates which channel: 0=x, 1=y, 2=z and 3=w.

Microcode VINTRP Opcode 0 (0x0)

1	1	0	1	0	1	VDST	OP	ATTR	ATTR CHAN	VSRC (I, J)	+0
---	---	---	---	---	---	------	----	------	--------------	-------------	----

Microcode VOP3a Opcode 624 (0x270)

NEG	OMOD	SRC2	SRC1	SRC0	+4						
1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0

Instruction **V_INTERP_P2_F32**

Description Vertex Parameter Interpolation using parameters stored in LDS and barycentric coordinates in VGPRs. Second step for interpolation: $D.f = P20.f * S0.f + D.f$.
M0 must contain: { 1'b0, new_prim_mask[15:1], lds_param_offset[15:0] }.
The ATTR field indicates which attribute (0-32) to interpolate.
The ATTRCHAN field indicates which channel: 0=x, 1=y, 2=z and 3=w.

Microcode VINTRP Opcode 1 (0x1)

1	1	0	1	0	1	VDST	OP	ATTR	ATTR CHAN	VSRC (I, J)	+0
---	---	---	---	---	---	------	----	------	--------------	-------------	----

Microcode VOP3a Opcode 625 (0x271)

NEG	OMOD	SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0

Instruction **V_INTERP_P1LL_F16**

Description 'LL' stands for 'two LDS arguments'. attr_word selects the high or low half 16 bits of each LDS dword accessed. This opcode is available for 32-bank LDS only. NOTE: In textual representations the I/J VGPR is the first source and the attribute is the second source; however in the VOP3 encoding the attribute is stored in the src0 field and the VGPR is stored in the src1 field.
 $D.f32 = P10.f16 * S0.f32 + P0.f16$.

Microcode VOP3a Opcode 628 (0x274)

NEG	OMOD	SRC2				SRC1			SRC0		+4
1	1	0	1	0	0	OP	CL MP	r	ABS	VDST	+0

Instruction **V_INTERP_P1LV_F16**

Description 'LV' stands for 'One LDS and one VGPR argument'. S2 holds two parameters, attr_word selects the high or low word of the VGPR for this calculation, as well as the high or low half of the LDS data. Meant for use with 16-bank LDS. NOTE: In textual representations the I/J VGPR is the first source and the attribute is the second source; however in the VOP3 encoding the attribute is stored in the src0 field and the VGPR is stored in the src1 field. encoding the attribute is stored in the src0 field and the VGPR is stored in the src1 field.
 $D.f32 = P10.f16 * S0.f32 + (S2.u32 \gg (attr_word * 16)).f16.$

Microcode VOP3a Opcode 629 (0x275)

NEG			OMOD			SRC2						SRC1				SRC0				+4
1	1	0	1	0	0	OP						CL	r	ABS	VDST				+0	
												MP								

Instruction **V_INTERP_P2_F16**

Description Final computation. attr_word selects LDS high or low 16bits. Used for both 16- and 32-bank LDS. Result is written to the 16 LSBs of the destination VGPR. NOTE: In textual representations, the I/J VGPR is the first source and the attribute is the second source; however, in the VOP3 encoding, the attribute is stored in the src0 field and the VGPR is stored in the src1 field.
 $D.f16 = P20.f16 * S0.f32 + S2.f32.$

Microcode VOP3a Opcode 630 (0x276)

NEG			OMOD			SRC2						SRC1				SRC0				+4
1	1	0	1	0	0	OP						CL	r	ABS	VDST				+0	
												MP								

12.13 LDS/GDS Instructions

This suite of instructions operates on data stored within the data share memory. The instructions transfer data between VGPRs and data share memory.

The bitfield map for the LDS/GDS is:

VDST							DATA1		DATA0		ADDR		+4	
1	1	0	1	1	0	r	OP		G D S	OFFSET1		OFFSET0		+0

where:

OFFSET0 = Unsigned byte offset added to the address supplied by the ADDR VGPR.

OFFSET1 = Unsigned byte offset added to the address supplied by the ADDR VGPR.

GDS = Set if GDS, cleared if LDS.

OP = DS instructions.

ADDR = Source LDS address VGPR 0 - 255.

DATA0 = Source data0 VGPR 0 - 255.

DATA1 = Source data1 VGPR 0 - 255.

VDST = Destination VGPR 0- 255.

Table 12.7 DS Instructions for the Opcode Field

Instruction	Description (C-Function Equivalent)	Decimal/Hex
DS_ADD_U32	DS[A] = DS[A] + D0; uint add.	00 (0x0)
DS_SUB_U32	DS[A] = DS[A] - D0; uint subtract.	01 (0x1)
DS_RSUB_U32	DS[A] = D0 - DS[A]; uint reverse subtract.	02 (0x2)
DS_INC_U32	DS[A] = (DS[A] >= D0 ? 0 : DS[A] + 1); uint increment.	03 (0x3)
DS_DEC_U32	DS[A] = (DS[A] == 0 DS[A] > D0 ? D0 : DS[A] - 1); uint decrement.	04 (0x4)
DS_MIN_I32	DS[A] = min(DS[A], D0); int min.	05 (0x5)
DS_MAX_I32	DS[A] = max(DS[A], D0); int max.	06 (0x6)
DS_MIN_U32	DS[A] = min(DS[A], D0); uint min.	07 (0x7)
DS_MAX_U32	DS[A] = max(DS[A], D0); uint max.	08 (0x8)
DS_AND_B32	DS[A] = DS[A] & D0; Dword AND.	09 (0x9)
DS_OR_B32	DS[A] = DS[A] D0; Dword OR.	10 (0xA)
DS_XOR_B32	DS[A] = DS[A] ^ D0; Dword XOR.	11 (0xB)
DS_MSKOR_B32	DS[A] = (DS[A] ^ ~D0) D1; masked Dword OR.	12 (0xC)
DS_WRITE_B32	DS[A] = D0; write a Dword.	13 (0xD)
DS_WRITE2_B32	DS[ADDR+offset0*4] = D0; DS[ADDR+offset1*4] = D1; write 2 Dwords.	14 (0xE)

Table 12.7 DS Instructions for the Opcode Field (Cont.)

Instruction	Description (C-Function Equivalent)	Decimal/Hex
DS_WRITE2ST64_B32	DS[ADDR+offset0*4*64] = D0; DS[ADDR+offset1*4*64] = D1; write 2 Dwords.	15 (0xF)
DS_CMPST_B32	DS[A] = (DS[A] == D0 ? D1 : DS[A]); compare store.	16 (0x10)
DS_CMPST_F32	DS[A] = (DS[A] == D0 ? D1 : DS[A]); compare store with float rules.	17 (0x11)
DS_MIN_F32	DS[A] = (DS[A] < D1) ? D0 : DS[A]; float compare swap (handles NaN/INF/denorm).	18 (0x12)
DS_MAX_F32	DS[A] = (D0 > DS[A]) ? D0 : DS[A]; float, handles NaN/INF/denorm.	19 (0x13)
DS_NOP	Do nothing.	20 (0x14)
DS_ADD_F32	DS[A] = DS[A] + D0; float add.	21 (0x15)
DS_WRITE_B8	DS[A] = D0[7:0]; byte write.	30 (0x1E)
DS_WRITE_B16	DS[A] = D0[15:0]; short write.	31 (0x1F)
DS_ADD_RTN_U32	Uint add.	32 (0x20)
DS_SUB_RTN_U32	Uint subtract.	33 (0x21)
DS_RSUB_RTN_U32	Uint reverse subtract.	34 (0x22)
DS_INC_RTN_U32	Uint increment.	35 (0x23)
DS_DEC_RTN_U32	Uint decrement.	36 (0x24)
DS_MIN_RTN_I32	Int min.	37 (0x25)
DS_MAX_RTN_I32	Int max.	38 (0x26)
DS_MIN_RTN_U32	Uint min.	39 (0x27)
DS_MAX_RTN_U32	Uint max.	40 (0x28)
DS_AND_RTN_B32	Dword AND.	41 (0x29)
DS_OR_RTN_B32	Dword OR.	42 (0x2A)
DS_XOR_RTN_B32	Dword XOR.	43 (0x2B)
DS_MSKOR_RTN_B32	Masked Dword OR.	44 (0x2C)
DS_WRXCHG_RTN_B32	Write exchange. Offset = {offset1,offset0}. A = ADDR+offset. D=DS[Addr]. DS[Addr]=D0.	45 (0x2D)
DS_WRXCHG2_RTN_B32	Write exchange 2 separate Dwords.	46 (0x2E)
DS_WRXCHG2ST64_RTN_B32	Write exchange 2 Dwords, stride 64.	47 (0x2F)
DS_CMPST_RTN_B32	Compare store.	48 (0x30)
DS_CMPST_RTN_F32	Compare store with float rules.	49 (0x31)
DS_MIN_RTN_F32	DS[A] = (DS[A] < D1) ? D0 : DS[A]; float compare swap (handles NaN/INF/denorm).	50 (0x32)
DS_MAX_RTN_F32	DS[A] = (D0 > DS[A]) ? D0 : DS[A]; float, handles NaN/INF/denorm.	51 (0x33)
DS_WRAP_RTN_B32	DS[A] = (DS[A] >= D0) ? DS[A] - D0 : DS[A] + D1.	52 (0x34)
DS_READ_B32	R = DS[A]; Dword read.	54 (0x36)
DS_READ2_B32	R = DS[ADDR+offset0*4], R+1 = DS[ADDR+offset1*4]. Read 2 Dwords.	55 (0x37)
DS_READ2ST64_B32	R = DS[ADDR+offset0*4*64], R+1 = DS[ADDR+offset1*4*64]. Read 2 Dwords.	56 (0x38)
DS_READ_I8	R = signext(DS[A][7:0]); signed byte read.	57 (0x39)
DS_READ_U8	R = {24'h0,DS[A][7:0]}; unsigned byte read.	58 (0x3A)
DS_READ_I16	R = signext(DS[A][15:0]); signed short read.	59 (0x3B)

Table 12.7 DS Instructions for the Opcode Field (Cont.)

Instruction	Description (C-Function Equivalent)	Decimal/Hex
DS_READ_U16	R = {16'h0,DS[A][15:0]}; unsigned short read.	60 (0x3C)
DS_SWIZZLE_B32	Swizzles input thread data based on offset mask and returns; note does not read or write the DS memory banks. <pre> offset = offset1:offset0; // full data sharing within 4 consecutive threads if (offset[15]) { for (i = 0; i < 32; i+=4) { thread_out[i+0] = thread_valid[i+offset[1:0]] ? thread_in[i+offset[1:0]] : 0; thread_out[i+1] = thread_valid[i+offset[3:2]] ? thread_in[i+offset[3:2]] : 0; thread_out[i+2] = thread_valid[i+offset[5:4]] ? thread_in[i+offset[5:4]] : 0; thread_out[i+3] = thread_valid[i+offset[7:6]] ? thread_in[i+offset[7:6]] : 0; } } // limited data sharing within 32 consecutive threads else { and_mask = offset[4:0]; or_mask = offset[9:5]; xor_mask = offset[14:10]; for (i = 0; i < 32; i++) { j = ((i & and_mask) or_mask) ^ xor_mask; thread_out[i] = thread_valid[j] ? thread_in[j] : 0; } } </pre>	61 (0x3D)
DS_PERMUTE_B32	Forward permute. Does not write any LDS memory. LDS[dst] = src0 returnVal = LDS[thread_id]. Where "thread_id" is 0..63.	62 (0x3E)
DS_PERMUTE_B32	Backward permute. Does not write any LDS memory. LDS[thread_id] = src0. Where "thread_id" is 0..63. returnVal = LDS[dst]	63 (0x3F)
DS_ADD_U64	Uint add.	64 (0x40)
DS_SUB_U64	Uint subtract.	65 (0x41)
DS_RSUB_U64	Uint reverse subtract.	66 (0x42)
DS_INC_U64	Uint increment.	67 (0x43)
DS_DEC_U64	Uint decrement.	68 (0x44)
DS_MIN_I64	Int min.	69 (0x45)
DS_MAX_I64	Int max.	70 (0x46)
DS_MIN_U64	Uint min.	71 (0x47)
DS_MAX_U64	Uint max.	72 (0x48)
DS_AND_B64	Dword AND.	73 (0x49)
DS_OR_B64	Dword OR.	74 (0x4A)
DS_XOR_B64	Dword XOR.	75 (0x4B)
DS_MSKOR_B64	Masked Dword XOR.	76 (0x4C)
DS_WRITE_B64	Write.	77 (0x4D)
DS_WRITE2_B64	DS[ADDR+offset0*8] = D0; DS[ADDR+offset1*8] = D1; write 2 Dwords.	78 (0x4E)
DS_WRITE2ST64_B64	DS[ADDR+offset0*8*64] = D0; DS[ADDR+offset1*8*64] = D1; write 2 Dwords.	79 (0x4F)
DS_CMPST_B64	Compare store.	80 (0x50)

Table 12.7 DS Instructions for the Opcode Field (Cont.)

Instruction	Description (C-Function Equivalent)	Decimal/Hex
DS_CMPST_F64	Compare store with float rules.	81 (0x51)
DS_MIN_F64	$DS[A] = (D0 < DS[A]) ? D0 : DS[A]$; float, handles NaN/INF/denorm.	82 (0x52)
DS_MAX_F64	$DS[A] = (D0 > DS[A]) ? D0 : DS[A]$; float, handles NaN/INF/denorm.	83 (0x53)
DS_ADD_RTN_U64	Uint add.	96 (0x60)
DS_SUB_RTN_U64	Uint subtract.	97 (0x61)
DS_RSUB_RTN_U64	Uint reverse subtract.	98 (0x62)
DS_INC_RTN_U64	Uint increment.	99 (0x63)
DS_DEC_RTN_U64	Uint decrement.	100 (0x64)
DS_MIN_RTN_I64	Int min.	101 (0x65)
DS_MAX_RTN_I64	Int max.	102 (0x66)
DS_MIN_RTN_U64	Uint min.	103 (0x67)
DS_MAX_RTN_U64	Uint max.	104 (0x68)
DS_AND_RTN_B64	Dword AND.	105 (0x69)
DS_OR_RTN_B64	Dword OR.	106 (0x6A)
DS_XOR_RTN_B64	Dword XOR.	107 (0x6B)
DS_MSKOR_RTN_B64	Masked Dword XOR.	108 (0x6C)
DS_WRXCHG_RTN_B64	Write exchange.	109 (0x6D)
DS_WRXCHG2_RTN_B64	Write exchange relative.	110 (0x6E)
DS_WRXCHG2ST64_RTN_B64	Write exchange 2 Dwords.	111 (0x6F)
DS_CMPST_RTN_B64	Compare store.	112 (0x70)
DS_CMPST_RTN_F64	Compare store with float rules.	113 (0x71)
DS_MIN_RTN_F64	$DS[A] = (D0 < DS[A]) ? D0 : DS[A]$; float, handles NaN/INF/denorm.	114 (0x72)
DS_MAX_RTN_F64	$DS[A] = (D0 > DS[A]) ? D0 : DS[A]$; float, handles NaN/INF/denorm.	115 (0x73)
DS_READ_B64	Dword read.	118 (0x74)
DS_READ2_B64	$R = DS[ADDR+offset0*8]$, $R+1 = DS[ADDR+offset1*8]$. Read 2 Dwords.	119 (0x75)
DS_READ2ST64_B64	$R = DS[ADDR+offset0*8*64]$, $R+1 = DS[ADDR+offset1*8*64]$. Read 2 Dwords.	120 (0x76)
DS_CONDXCHG32_RTN_B64	Conditional write exchange.	126 (0x7E)
DS_ADD_SRC2_U32	$B = A + 4*(offset1[7] ? \{A[31],A[31:17]\} : \{offset1[6],offset1[6:0],offset0\})$. $DS[A] = DS[A] + DS[B]$; uint add.	128 (0x80)
DS_SUB_SRC2_U32	$B = A + 4*(offset1[7] ? \{A[31],A[31:17]\} : \{offset1[6],offset1[6:0],offset0\})$. $DS[A] = DS[A] - DS[B]$; uint subtract.	129 (0x81)
DS_RSUB_SRC2_U32	$B = A + 4*(offset1[7] ? \{A[31],A[31:17]\} : \{offset1[6],offset1[6:0],offset0\})$. $DS[A] = DS[B] - DS[A]$; uint reverse subtract.	130 (0x82)

Table 12.7 DS Instructions for the Opcode Field (Cont.)


Instruction	Description (C-Function Equivalent)	Decimal/Hex
DS_INC_SRC2_U32	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = (DS[A] >= DS[B] ? 0 : DS[A] + 1); uint increment.	131 (0x83)
DS_DEC_SRC2_U32	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = (DS[A] == 0 DS[A] > DS[B] ? DS[B] : DS[A] - 1); uint decrement.	132 (0x84)
DS_MIN_SRC2_I32	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = min(DS[A], DS[B]); int min.	133 (0x85)
DS_MAX_SRC2_I32	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = max(DS[A], DS[B]); int max.	134 (0x86)
DS_MIN_SRC2_U32	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = min(DS[A], DS[B]); uint min.	135 (0x87)
DS_MAX_SRC2_U32	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = max(DS[A], DS[B]); uint maxw	136 (0x88)
DS_AND_SRC2_B32	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = DS[A] & DS[B]; Dword AND.	137 (0x89)
DS_OR_SRC2_B32	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = DS[A] DS[B]; Dword OR.	138 (0x8A)
DS_XOR_SRC2_B32	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = DS[A] ^ DS[B]; Dword XOR.	139 (0x8B)
DS_WRITE_SRC2_B32	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = DS[B]; write Dword.	141 (0x8D)
DS_MIN_SRC2_F32	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = (DS[B] < DS[A]) ? DS[B] : DS[A]; float, handles NaN/INF/denorm.	146 (0x92)
DS_MAX_SRC2_F32	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = (DS[B] > DS[A]) ? DS[B] : DS[A]; float, handles NaN/INF/denorm.	147 (0x93)
DS_GWS_SEMA_RELEASE_ALL	GDS Only. Release all wavefronts waiting on this semaphore. ResourceID is in offset[4:0].	152 (0x98)
DS_GWS_INIT	GDS only.	153 (0x99)
DS_GWS_SEMA_V	GDS only.	154 (0x9A)
DS_GWS_SEMA_BR	GDS only.	155 (0x9B)
DS_GWS_SEMA_P	GDS only.	156 (0x9C)
DS_GWS_BARRIER	GDS only.	157 (0x9D)
DS_CONSUME	Consume entries from a buffer.	189 (0xBD)
DS_APPEND	Append one or more entries to a buffer.	190 (0xBE)
DS_ORDERED_COUNT	Increment an append counter. The operation is done in wavefront-creation order.	191 (0xBF)
DS_ADD_SRC2_U64	Uint add.	192 (0xC0)

Table 12.7 DS Instructions for the Opcode Field (Cont.)

Instruction	Description (C-Function Equivalent)	Decimal/Hex
DS_SUB_SRC2_U64	Uint subtract.	193 (0xC1)
DS_RSUB_SRC2_U64	Uint reverse subtract.	194 (0xC2)
DS_INC_SRC2_U64	Uint increment.	195 (0xC3)
DS_DEC_SRC2_U64	Uint decrement.	196 (0xC4)
DS_MIN_SRC2_I64	Int min.	197 (0xC5)
DS_MAX_SRC2_I64	Int max.	198 (0xC6)
DS_MIN_SRC2_U64	Uint min.	199 (0xC7)
DS_MAX_SRC2_U64	Uint max.	200 (0xC8)
DS_AND_SRC2_B64	Dword AND.	201 (0xC9)
DS_OR_SRC2_B64	Dword OR.	202 (0xCA)
DS_XOR_SRC2_B64	Dword XOR.	203 (0xCB)
DS_WRITE_SRC2_B64	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = DS[B]; write Qword.	205 (0xCD)
DS_MIN_SRC2_F64	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. [A] = (D0 < DS[A]) ? D0 : DS[A]; float, handles NaN/INF/denorm.	210 (0xD2)
DS_MAX_SRC2_F64	$B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. [A] = (D0 > DS[A]) ? D0 : DS[A]; float, handles NaN/INF/denorm.	211 (0xD3)
DS_WRITE_B96	{DS[A+2], DS[A+1], DS[A]} = D0[95:0]; tri-dword write.	222 (0xDE)
DS_WRITE_B128	{DS[A+3], DS[A+2], DS[A+1], DS[A]} = D0[127:0]; qword write.	223 (0xDF)
DS_CONDXCHG32_RTN_B128	Conditional write exchange.	253 (0xFD)
DS_READ_B96	Tri-dword read.	254 (0xFE)
DS_READ_B128	Qword read.	255 (0xFF)

12.14 MUBUF Instructions

The bitfield map of the MUBUF format is:



where:

- OFFSET = Unsigned byte offset.
- OFFEN = Send offset either as VADDR or as zero..
- IDXEN = Send index either as VADDR or as zero.
- GLC = Global coherency.
- LDS = Data read from/written to LDS or VGPR.
- OP = Opcode instructions.
- VADDR = VGPR address source.
- VDATA = Destination vector GPR.
- SRSRC = Scalar GPR that specifies resource constant.
- SLC = System level coherent.
- TFE = Texture fail enable.
- SOFFSET = Byte offset added to the memory address.

Table 12.8 MUBUF Instructions for the Opcode Field

Instruction	Description	Decimal (Hex)
LOAD FORMAT		
BUFFER_LOAD_FORMAT_X	Untyped buffer load 1 Dword with format conversion.	0 (0x0)
BUFFER_LOAD_FORMAT_XY	Untyped buffer load 2 Dwords with format conversion.	1 (0x1)
BUFFER_LOAD_FORMAT_XYZ	Untyped buffer load 3 Dwords with format conversion.	2 (0x2)
BUFFER_LOAD_FORMAT_XYZW	Untyped buffer load 4 Dwords with format conversion.	3 (0x3)
STORE FORMAT		
BUFFER_STORE_FORMAT_X	Untyped buffer store 1 Dword with format conversion.	4 (0x4)
BUFFER_STORE_FORMAT_XY	Untyped buffer store 2 Dwords with format conversion.	5 (0x5)
BUFFER_STORE_FORMAT_XYZ	Untyped buffer store 3 Dwords with format conversion.	6 (0x6)
BUFFER_STORE_FORMAT_XYZW	Untyped buffer store 4 Dwords with format conversion.	7 (0x7)
LOAD		
BUFFER_LOAD_FORMAT_D16_X	Untyped buffer load 1 dword with format conversion.	8 (0x8)
BUFFER_LOAD_FORMAT_D16_XY	Untyped buffer load 2 dwords with format conversion.	9 (0x9)
BUFFER_LOAD_FORMAT_D16_XYZ	Untyped buffer load 3 dwords with format conversion.	10 (0xA)

Table 12.8 MUBUF Instructions for the Opcode Field (Cont.)


Instruction	Description	Decimal (Hex)
BUFFER_LOAD_FORMAT_D16_XYZW	Untyped buffer load 4 dwords with format conversion.	11 (0xB)
BUFFER_LOAD_UBYTE	Untyped buffer load unsigned byte.	16 (0x10)
BUFFER_LOAD_SBYTE	Untyped buffer load signed byte.	17 (0x11)
BUFFER_LOAD_USHORT	Untyped buffer load unsigned short.	18 (0x12)
BUFFER_LOAD_SSHORT	Untyped buffer load signed short.	19 (0x13)
BUFFER_LOAD_DWORD	Untyped buffer load Dword.	20 (0x14)
BUFFER_LOAD_DWORDX2	Untyped buffer load 2 Dwords.	21 (0x15)
BUFFER_LOAD_DWORDX3	Untyped buffer load 3 Dwords.	22 (0x16)
BUFFER_LOAD_DWORDX4	Untyped buffer load 4 Dwords.	23 (0x17)
STORE		
BUFFER_STORE_FORMAT_D16_X	Untyped buffer store 1 dword with format conversion.	12 (0xC)
BUFFER_STORE_FORMAT_D16_XY	Untyped buffer load 2 dwords with format conversion.	13 (0xD)
BUFFER_STORE_FORMAT_D16_XYZ	Untyped buffer load 3 dwords with format conversion.	14 (0xE)
BUFFER_STORE_FORMAT_D16_XYZW	Untyped buffer load 4 dwords with format conversion.	15 (0xF)
BUFFER_STORE_BYTE	Untyped buffer store byte.	24 (0x18)
BUFFER_STORE_SHORT	Untyped buffer store short.	26 (0x1A)
BUFFER_STORE_DWORD	Untyped buffer store Dword.	28 (0x1C)
BUFFER_STORE_DWORDX2	Untyped buffer store 2 Dwords.	29 (0x1D)
BUFFER_STORE_DWORDX3	Untyped buffer store 3 Dwords.	30 (0x1E)
BUFFER_STORE_DWORDX4	Untyped buffer store 4 Dwords.	31 (0x1F)
BUFFER_STORE_LDS_DWORD	Store one Dword from LDS memory to system memory without using VGPRs.	61 (0x3D)
Cache Invalidation		
BUFFER_WBINVL1	Write back and invalidate the shader L1. Always returns ACK to shader.	62 (0x2E)
BUFFER_WBINVL1_VOL	Write back and invalidate the shader L1 only for lines of MTYPE SC and GC. Always returns ACK to shader.	63 (0x3F)
ATOMIC		
BUFFER_ATOMIC_SWAP	32b, dst=src, returns previous value if glc==1.	64 (0x40)
BUFFER_ATOMIC_CMPSWAP	32b, dst = (dst==cmp) ? src : dst. Returns previous value if glc==1. src comes from the first data-vgpr, cmp from the second.	65 (0x41)
BUFFER_ATOMIC_ADD	32b, dst += src. Returns previous value if glc==1.	66 (0x42)
BUFFER_ATOMIC_SUB	32b, dst -= src. Returns previous value if glc==1.	67 (0x43)
BUFFER_ATOMIC_SMIN	32b, dst = (src < dst) ? src : dst (signed). Returns previous value if glc==1.	68 (0x44)
BUFFER_ATOMIC_UMIN	32b, dst = (src < dst) ? src : dst (unsigned). Returns previous value if glc==1.	69 (0x45)
BUFFER_ATOMIC_SMAX	32b, dst = (src > dst) ? src : dst (signed). Returns previous value if glc==1.	70 (0x46)

Table 12.8 MUBUF Instructions for the Opcode Field (Cont.)

Instruction	Description	Decimal (Hex)
BUFFER_ATOMIC_UMAX	32b, dst = (src > dst) ? src : dst (unsigned). Returns previous value if glc==1.	71 (0x47)
BUFFER_ATOMIC_AND	32b, dst &= src. Returns previous value if glc==1.	72 (0x48)
BUFFER_ATOMIC_OR	32b, dst = src. Returns previous value if glc==1.	73 (0x49)
BUFFER_ATOMIC_XOR	32b, dst ^= src. Returns previous value if glc==1.	74 (0x4A)
BUFFER_ATOMIC_INC	32b, dst = (dst >= src) ? 0 : dst+1. Returns previous value if glc==1.	75 (0x4B)
BUFFER_ATOMIC_DEC	32b, dst = ((dst==0 (dst > src)) ? src : dst-1). Returns previous value if glc==1.	76 (0x4C)
BUFFER_ATOMIC_SWAP_X2	64b, dst=src, returns previous value if glc==1.	96 (0x60)
BUFFER_ATOMIC_CMPSWAP_X2	64b, dst = (dst==cmp) ? src : dst. Returns previous value if glc==1. src comes from the first two data-vgprs, cmp from the second two.	97 (0x61)
BUFFER_ATOMIC_ADD_X2	64b, dst += src. Returns previous value if glc==1.	98 (0x62)
BUFFER_ATOMIC_SUB_X2	64b, dst -= src. Returns previous value if glc==1.	99 (0x63)
BUFFER_ATOMIC_SMIN_X2	64b, dst = (src < dst) ? src : dst (signed). Returns previous value if glc==1.	100 (0x64)
BUFFER_ATOMIC_UMIN_X2	64b, dst = (src < dst) ? src : dst (unsigned). Returns previous value if glc==1.	101 (0x65)
BUFFER_ATOMIC_SMAX_X2	64b, dst = (src > dst) ? src : dst (signed). Returns previous value if glc==1.	102 (0x66)
BUFFER_ATOMIC_UMAX_X2	64b, dst = (src > dst) ? src : dst (unsigned). Returns previous value if glc==1.	103 (0x67)
BUFFER_ATOMIC_AND_X2	64b, dst &= src. Returns previous value if glc==1.	104 (0x68)
BUFFER_ATOMIC_OR_X2	64b, dst = src. Returns previous value if glc==1.	105 (0x69)
BUFFER_ATOMIC_XOR_X2	64b, dst ^= src. Returns previous value if glc==1.	106 (0x6A)
BUFFER_ATOMIC_INC_X2	64b, dst = (dst >= src) ? 0 : dst+1. Returns previous value if glc==1.	107 (0x6B)
BUFFER_ATOMIC_DEC_X2	64b, dst = ((dst==0 (dst > src)) ? src : dst-1). Returns previous value if glc==1.	108 (0x6C)

12.15 MTBUF Instructions

The bitfield map of the MTBUF format is:



where:

- OFFSET = Unsigned byte offset.
- OFFEN = Send offset either as VADDR or as zero.
- IDXEN = Send index either as VADDR or as zero.
- GLC = Global coherency.
- OP = Opcode instructions.
- DFMT = Data format for typed buffer.
- NFMT = Number format for typed buffer.
- VADDR = VGPR address source.
- VDATA = Vector GPR for read/write result.
- SRSRC = Scalar GPR that specifies resource constant.

Table 12.9 MTBUF Instructions for the Opcode Field

Instruction	Description	Decimal (Hex)
LOAD		
TBUFFER_LOAD_FORMAT_X	Typed buffer load 1 Dword with format conversion.	0 (0x0)
TBUFFER_LOAD_FORMAT_XY	Typed buffer load 2 Dwords with format conversion.	1 (0x1)
TBUFFER_LOAD_FORMAT_XYZ	Typed buffer load 3 Dwords with format conversion.	2 (0x2)
TBUFFER_LOAD_FORMAT_XYZW	Typed buffer load 4 Dwords with format conversion.	3 (0x3)
STORE		
TBUFFER_STORE_FORMAT_X	Typed buffer store 1 Dword with format conversion.	4 (0x4)
TBUFFER_STORE_FORMAT_XY	Typed buffer store 2 Dwords with format conversion.	5 (0x5)
TBUFFER_STORE_FORMAT_XYZ	Typed buffer store 3 Dwords with format conversion.	6 (0x6)
TBUFFER_STORE_FORMAT_XYZW	Typed buffer store 4 Dwords with format conversion.	7 (0x7)
TBUFFER_LOAD_FORMAT_D16_X	Typed buffer load 1 Dword with format conversion.	8 (0x8)
TBUFFER_LOAD_FORMAT_D16_XY	Typed buffer load 2 Dwords with format conversion.	9 (0x9)
TBUFFER_LOAD_FORMAT_D16_XYZ	Typed buffer load 3 Dwords with format conversion.	10 (0xA)
TBUFFER_LOAD_FORMAT_D16_XYZW	Typed buffer load 4 Dwords with format conversion.	11 (0xB)
TBUFFER_STORE_FORMAT_D16_X	Typed buffer store 1 Dword with format conversion.	12 (0xC)

Table 12.9 MTBUF Instructions for the Opcode Field (Cont.)

Instruction	Description	Decimal (Hex)
TBUFFER_STORE_FORMAT_D16_XY	Typed buffer store 2 Dwords with format conversion.	13 (0xD)
TBUFFER_STORE_FORMAT_D16_XYZ	Typed buffer store 3 Dwords with format conversion.	14 (0xE)
TBUFFER_STORE_FORMAT_D16_XYZW	Typed buffer store 4 Dwords with format conversion.	15 (0xF)

Table 12.10 NFMT: Shader Num_Format

Value	Encode	Buffer r	Buffer w
0	unorm	yes	yes
1	snorm	yes	yes
2	uscaled	yes	no
3	sscaled	yes	no
4	uint	yes	yes
5	sint	yes	yes
6	reserved		
7	float	yes	yes
8	reserved		
9	srgb	no	no
10-15	reserved		

Table 12.11 DFMT: Data_Format

Value	Encode
0	invalid
1	8
2	16
3	8_8
4	32
5	16_16
6	10_11_11
7	11_11_10

Value	Encode
8	10_10_10_2
9	2_10_10_10
10	8_8_8_8
11	32_32
12	16_16_16_16
13	32_32_32
14	32_32_32_32
15	reserved

12.16 MIMG Instructions

The bitfield map of the MTBUF format is:

D 16	r					SSAMP		SRSRC			VDATA				VADDR		+4	
1	1	1	1	0	0	SLC	OP			LWE	TFE	R128	DA	GLC	UNRM	DMASK	r	+0

where:

- DMASK = Enable mask for image read/write data components.
- UNRM = Force address to be unnormalized.
- GLC = Global coherency.
- DA = Declare an array.
- R128 = Texture resource size.
- TFE = Texture fail enable.
- LWE = LOD warning enable.
- OP = Opcode instructions.
- SLC = System level coherent.
- VADDR = VGPR address source.
- VDATA = Vector GPR for read/write result.
- SRSRC = Scalar GPR that specifies resource constant.
- SSAMP = Scalar GPR that specifies sampler constant.

Table 12.12 MIMG Instructions for the Opcode Field

Instruction	Description	Decimal (Hex)
LOAD		
IMAGE_LOAD	Image memory load with format conversion specified in T#. No sampler.	0 (0x0)
IMAGE_LOAD_MIP	Image memory load with user-supplied mip level. No sampler.	1 (0x1)
IMAGE_LOAD_PCK	Image memory load with no format conversion. No sampler.	2 (0x2)
IMAGE_LOAD_PCK_SGN	Image memory load with no format conversion and sign extension. No sampler.	3 (0x3)
IMAGE_LOAD_MIP_PCK	Image memory load with user-supplied mip level, no format conversion. No sampler.	4 (0x4)
IMAGE_LOAD_MIP_PCK_SGN	Image memory load with user-supplied mip level, no format conversion and with sign extension. No sampler.	5 (0x5)
STORE		
IMAGE_STORE	Image memory store with format conversion specified in T#. No sampler.	8 (0x8)
IMAGE_STORE_MIP	Image memory store with format conversion specified in T# to user specified mip level. No sampler.	9 (0x9)

Table 12.12 MIMG Instructions for the Opcode Field (Cont.)

Instruction	Description	Decimal (Hex)
IMAGE_STORE_PCK	Image memory store of packed data without format conversion. No sampler.	10 (0xA)
IMAGE_STORE_MIP_PCK	Image memory store of packed data without format conversion to user-supplied mip level. No sampler.	11 (0xB)
ATOMIC		
IMAGE_ATOMIC_SWAP	dst=src, returns previous value if glc==1.	15 (0xF)
IMAGE_ATOMIC_CMPSWAP	dst = (dst==cmp) ? src : dst. Returns previous value if glc==1.	16 (0x10)
IMAGE_ATOMIC_ADD	dst += src. Returns previous value if glc==1.	17 (0x11)
IMAGE_ATOMIC_SUB	dst -= src. Returns previous value if glc==1.	18 (0x12)
IMAGE_ATOMIC_SMIN	dst = (src < dst) ? src : dst (signed). Returns previous value if glc==1.	20 (0x14)
IMAGE_ATOMIC_UMIN	dst = (src < dst) ? src : dst (unsigned). Returns previous value if glc==1.	21 (0x15)
IMAGE_ATOMIC_SMAX	dst = (src > dst) ? src : dst (signed). Returns previous value if glc==1.	22 (0x16)
IMAGE_ATOMIC_UMAX	dst = (src > dst) ? src : dst (unsigned). Returns previous value if glc==1.	23 (0x17)
IMAGE_ATOMIC_AND	dst &= src. Returns previous value if glc==1.	24 (0x18)
IMAGE_ATOMIC_OR	dst = src. Returns previous value if glc==1.	25 (0x19)
IMAGE_ATOMIC_XOR	dst ^= src. Returns previous value if glc==1.	26 (0x1A)
IMAGE_ATOMIC_INC	dst = (dst >= src) ? 0 : dst+1. Returns previous value if glc==1.	27 (0x1B)
IMAGE_ATOMIC_DEC	dst = ((dst==0 (dst > src)) ? src : dst)-1. Returns previous value if glc==1.	28 (0x1C)
SAMPLE		
IMAGE_SAMPLE	Sample texture map.	32 (0x20)
IMAGE_SAMPLE_CL	Sample texture map, with LOD clamp specified in shader.	33 (0x21)
IMAGE_SAMPLE_D	Sample texture map, with user derivatives.	34 (0x22)
IMAGE_SAMPLE_D_CL	Sample texture map, with LOD clamp specified in shader, with user derivatives.	35 (0x23)
IMAGE_SAMPLE_L	Sample texture map, with user LOD.	36 (0x24)
IMAGE_SAMPLE_B	Sample texture map, with lod bias.	37 (0x25)
IMAGE_SAMPLE_B_CL	Sample texture map, with LOD clamp specified in shader, with lod bias.	38 (0x26)
IMAGE_SAMPLE_LZ	Sample texture map, from level 0.	39 (0x27)
IMAGE_SAMPLE_C	Sample texture map, with PCF.	40 (0x28)
IMAGE_SAMPLE_C_CL	SAMPLE_C, with LOD clamp specified in shader.	41 (0x29)
IMAGE_SAMPLE_C_D	SAMPLE_C, with user derivatives.	42 (0x2A)
IMAGE_SAMPLE_C_D_CL	SAMPLE_C, with LOD clamp specified in shader, with user derivatives.	43 (0x2B)
IMAGE_SAMPLE_C_L	SAMPLE_C, with user LOD.	44 (0x2C)
IMAGE_SAMPLE_C_B	SAMPLE_C, with lod bias.	45 (0x2D)
IMAGE_SAMPLE_C_B_CL	SAMPLE_C, with LOD clamp specified in shader, with lod bias.	46 (0x2E)

Table 12.12 MIMG Instructions for the Opcode Field (Cont.)

Instruction	Description	Decimal (Hex)
IMAGE_SAMPLE_C_LZ	SAMPLE_C, from level 0.	47 (0x2F)
IMAGE_SAMPLE_O	Sample texture map, with user offsets.	48 (0x30)
IMAGE_SAMPLE_CL_O	SAMPLE_O with LOD clamp specified in shader.	49 (0x31)
IMAGE_SAMPLE_D_O	SAMPLE_O, with user derivatives.	50 (0x32)
IMAGE_SAMPLE_D_CL_O	SAMPLE_O, with LOD clamp specified in shader, with user derivatives.	51 (0x33)
IMAGE_SAMPLE_L_O	SAMPLE_O, with user LOD.	52 (0x34)
IMAGE_SAMPLE_B_O	SAMPLE_O, with lod bias.	53 (0x35)
IMAGE_SAMPLE_B_CL_O	SAMPLE_O, with LOD clamp specified in shader, with lod bias.	54 (0x36)
IMAGE_SAMPLE_LZ_O	SAMPLE_O, from level 0.	55 (0x37)
IMAGE_SAMPLE_C_O	SAMPLE_C with user specified offsets.	56 (0x38)
IMAGE_SAMPLE_C_CL_O	SAMPLE_C_O, with LOD clamp specified in shader.	57 (0x39)
IMAGE_SAMPLE_C_D_O	SAMPLE_C_O, with user derivatives.	58 (0x3A)
IMAGE_SAMPLE_C_D_CL_O	SAMPLE_C_O, with LOD clamp specified in shader, with user derivatives.	59 (0x3B)
IMAGE_SAMPLE_C_L_O	SAMPLE_C_O, with user LOD.	60 (0x3C)
IMAGE_SAMPLE_C_B_O	SAMPLE_C_O, with lod bias.	61 (0x3D)
IMAGE_SAMPLE_C_B_CL_O	SAMPLE_C_O, with LOD clamp specified in shader, with lod bias.	62 (0x3E)
IMAGE_SAMPLE_C_LZ_O	SAMPLE_C_O, from level 0.	63 (0x3F)
IMAGE_SAMPLE_CD	Sample texture map, with user derivatives (LOD per quad).	104 (0x68)
IMAGE_SAMPLE_CD_CL	Sample texture map, with LOD clamp specified in shader, with user derivatives (LOD per quad).	105 (0x69)
IMAGE_SAMPLE_C_CD	SAMPLE_C, with user derivatives (LOD per quad).	106 (0x6A)
IMAGE_SAMPLE_C_CD_CL	SAMPLE_C, with LOD clamp specified in shader, with user derivatives (LOD per quad).	107 (0x6B)
IMAGE_SAMPLE_CD_O	SAMPLE_O, with user derivatives (LOD per quad).	108 (0x6C)
IMAGE_SAMPLE_CD_CL_O	SAMPLE_O, with LOD clamp specified in shader, with user derivatives (LOD per quad).	109 (0x6D)
IMAGE_SAMPLE_C_CD_O	SAMPLE_C_O, with user derivatives (LOD per quad).	110 (0x6E)
IMAGE_SAMPLE_C_CD_CL_O	SAMPLE_C_O, with LOD clamp specified in shader, with user derivatives (LOD per quad).	111 (0x6F)
GATHER4		
IMAGE_GATHER4	gather 4 single component elements (2x2).	64 (0x40)
IMAGE_GATHER4_CL	gather 4 single component elements (2x2) with user LOD clamp.	65 (0x41)
IMAGE_GATHER4_L	gather 4 single component elements (2x2) with user LOD.	66 (0x42)
IMAGE_GATHER4_B	gather 4 single component elements (2x2) with user bias.	67 (0x43)
IMAGE_GATHER4_B_CL	gather 4 single component elements (2x2) with user bias and clamp.	68 (0x44)
IMAGE_GATHER4_LZ	gather 4 single component elements (2x2) at level 0.	69 (0x45)
IMAGE_GATHER4_C	gather 4 single component elements (2x2) with PCF.	70 (0x46)
IMAGE_GATHER4_C_CL	gather 4 single component elements (2x2) with user LOD clamp and PCF.	71 (0x47)

Table 12.12 MIMG Instructions for the Opcode Field (Cont.)

Instruction	Description	Decimal (Hex)
IMAGE_GATHER4_C_L	gather 4 single component elements (2x2) with user LOD and PCF.	76 (0x4C)
IMAGE_GATHER4_C_B	gather 4 single component elements (2x2) with user bias and PCF.	77 (0x4D)
IMAGE_GATHER4_C_B_CL	gather 4 single component elements (2x2) with user bias, clamp and PCF.	78 (0x4E)
IMAGE_GATHER4_C_LZ	gather 4 single component elements (2x2) at level 0, with PCF.	79 (0x4F)
IMAGE_GATHER4_O	GATHER4, with user offsets.	80 (0x50)
IMAGE_GATHER4_CL_O	GATHER4_CL, with user offsets.	81 (0x51)
IMAGE_GATHER4_L_O	GATHER4_L, with user offsets.	84 (0x54)
IMAGE_GATHER4_B_O	GATHER4_B, with user offsets.	85 (0x55)
IMAGE_GATHER4_B_CL_O	GATHER4_B_CL, with user offsets.	86 (0x56)
IMAGE_GATHER4_LZ_O	GATHER4_LZ, with user offsets.	87 (0x57)
IMAGE_GATHER4_C_O	GATHER4_C, with user offsets.	88 (0x58)
IMAGE_GATHER4_C_CL_O	GATHER4_C_CL, with user offsets.	89 (0x59)
IMAGE_GATHER4_C_L_O	GATHER4_C_L, with user offsets.	92 (0x5C)
IMAGE_GATHER4_C_B_O	GATHER4_B, with user offsets.	93 (0x5D)
IMAGE_GATHER4_C_B_CL_O	GATHER4_B_CL, with user offsets.	94 (0x5E)
IMAGE_GATHER4_C_LZ_O	GATHER4_C_LZ, with user offsets.	95 (0x5F)
Miscellaneous		
IMAGE_GET_RESINFO	No sampler. Returns resource info into four VGPRs for the specified MIP level. These are 32-bit integer values: Vdata3-0 = { #mipLevels, depth, height, width } For cubemaps, depth = 6 * Number_of_array_slices.	14 (0xE)
IMAGE_GET_LOD	Return calculated LOD.	96 (0x60)

12.17 EXP Instructions

Instruction **EXPORT**


Description Transfer vertex position, vertex parameter, pixel color, or pixel depth information to the output buffer.
 Every pixel shader must do at least one export to a color, depth or NULL target with the VM bit set to 1. This communicates the pixel-valid mask to the color and depth buffers. Every pixel does only one of the above export types with the DONE bit set to 1.
 Vertex shaders must do one or more position exports, and at least one parameter export. The final position export must have the DONE bit set to 1.

Microcode EXP

VSRC3						VSRC2						VSRC1			VSRC0			+4	
1	1	0	0	0	1	r						V	DO	C	TARGET		EN		+0
						M	NE	OMP											

12.18 FLAT Instructions

The bitfield map of the FLAT format is:



where:

- GLC = Global coherency.
- SLC = System level coherency.
- OP = Opcode instructions.
- ADDR = Source of flat address VGPR.
- DATA = Source data.
- TFE = Texture fail enable.
- VDST = Destination VGPR.

Table 12.13 FLAT Instructions for the Opcode Field

Instruction	Description	Decimal (Hex)
LOAD		
FLAT_LOAD_UBYTE	Flat load unsigned byte. Zero extend to VGPR destination.	8 (0x8)
FLAT_LOAD_SBYTE	Flat load signed byte. Sign extend to VGPR destination.	9 (0x9)
FLAT_LOAD_USHORT	Flat load unsigned short. Zero extend to VGPR destination.	10 (0xA)
FLAT_LOAD_SSHORT	Flat load signed short. Sign extend to VGPR destination.	11 (0xB)
FLAT_LOAD_DWORD	Flat load Dword.	12 (0xC)
FLAT_LOAD_DWORDX2	Flat load 2 Dwords.	13 (0xD)
FLAT_LOAD_DWORDX4	Flat load 4 Dwords.	14 (0xE)
FLAT_LOAD_DWORDX3	Flat load 3 Dwords.	15 (0xF)
STORE		
FLAT_STORE_BYTE	Flat store byte.	24 (0x18)
FLAT_STORE_SHORT	Flat store short.	26 (0x1A)
FLAT_STORE_DWORD	Flat store Dword.	28 (0x1C)
FLAT_STORE_DWORDX2	Flat store 2 Dwords.	29 (0x1D)
FLAT_STORE_DWORDX4	Flat store 4 Dwords.	30 (0x1E)
FLAT_STORE_DWORDX3	Flat store 3 Dwords.	31 (0x1F)
ATOMIC		
FLAT_ATOMIC_SWAP	32b. dst=src, returns previous value if rtn==1.	48 (0x30)

Table 12.13 FLAT Instructions for the Opcode Field (Cont.)

Instruction	Description	Decimal (Hex)
FLAT_ATOMIC_CMPSWAP	32b, dst = (dst==cmp) ? src : dst. Returns previous value if rtn==1. src comes from the first data-VGPR, cmp from the second.	49 (0x31)
FLAT_ATOMIC_ADD	32b, dst += src. Returns previous value if rtn==1.	50 (0x32)
FLAT_ATOMIC_SUB	32b, dst -= src. Returns previous value if rtn==1.	51 (0x33)
FLAT_ATOMIC_SMIN	32b, dst = (src < dst) ? src : dst (signed comparison). Returns previous value if rtn==1.	53 (0x35)
FLAT_ATOMIC_UMIN	32b, dst = (src < dst) ? src : dst (unsigned comparison). Returns previous value if rtn==1.	54 (0x36)
FLAT_ATOMIC_SMAX	32b, dst = (src > dst) ? src : dst (signed comparison). Returns previous value if rtn==1.	55 (0x37)
FLAT_ATOMIC_UMAX	32b, dst = (src > dst) ? src : dst (unsigned comparison). Returns previous value if rtn==1.	56 (0x38)
FLAT_ATOMIC_AND	32b, dst &= src. Returns previous value if rtn==1.	57 (0x39)
FLAT_ATOMIC_OR	32b, dst = src. Returns previous value if rtn==1.	58 (0x3A)
FLAT_ATOMIC_XOR	32b, dst ^= src. Returns previous value if rtn==1.	59 (0x3B)
FLAT_ATOMIC_INC	32b, dst = (dst >= src) ? 0 : dst+1 (unsigned comparison). Returns previous value if rtn==1.	60 (0x3C)
FLAT_ATOMIC_DEC	32b, dst = ((dst==0 (dst > src)) ? src : dst-1 (unsigned comparison). Returns previous value if rtn==1.	61 (0x3D)
FLAT_ATOMIC_CMPSWAP_X2	64b, dst = (dst==cmp) ? src : dst. Returns previous value if rtn==1. src comes from the first two data-VGPRs, cmp from the second two.	81 (0x51)
FLAT_ATOMIC_ADD_X2	64b, dst += src. Returns previous value if rtn==1.	82 (0x52)
FLAT_ATOMIC_SUB_X2	64b, dst -= src. Returns previous value if rtn==1.	83 (0x53)
FLAT_ATOMIC_SMIN_X2	64b, dst = (src < dst) ? src : dst (signed comparison). Returns previous value if rtn==1.	85 (0x55)
FLAT_ATOMIC_UMIN_X2	64b, dst = (src < dst) ? src : dst (unsigned comparison). Returns previous value if rtn==1.	86 (0x56)
FLAT_ATOMIC_SMAX_X2	64b, dst = (src > dst) ? src : dst (signed comparison). Returns previous value if rtn==1.	87 (0x57)
FLAT_ATOMIC_UMAX_X2	64b, dst = (src > dst) ? src : dst (unsigned comparison). Returns previous value if rtn==1.	88 (0x58)
FLAT_ATOMIC_AND_X2	64b, dst &= src. Returns previous value if rtn==1.	89 (0x59)
FLAT_ATOMIC_OR_X2	64b, dst = src. Returns previous value if rtn==1.	90 (0x5A)
FLAT_ATOMIC_XOR_X2	64b, dst ^= src. Returns previous value if rtn==1.	91 (0x5B)
FLAT_ATOMIC_INC_X2	64b, dst = (dst >= src) ? 0 : dst+1. Returns previous value if rtn==1.	92 (0x5C)
FLAT_ATOMIC_DEC_X2	64b, dst = ((dst==0 (dst > src)) ? src : dst - 1. Returns previous value if rtn==1.	93 (0x5D)

Chapter 13

Microcode Formats

This section specifies the microcode formats. The definitions can be used to simplify compilation by providing standard templates and enumeration names for the various instruction formats.

Endian Order – The GCN architecture addresses memory and registers using little-endian byte-ordering and bit-ordering. Multi-byte values are stored with their least-significant (low-order) byte (LSB) at the lowest byte address, and they are illustrated with their LSB at the right side. Byte values are stored with their least-significant (low-order) bit (lsb) at the lowest bit address, and they are illustrated with their lsb at the right side.

Table 13.1 summarizes the microcode formats and their widths. The sections that follow provide details.

Table 13.1 Summary of Microcode Formats

Microcode Formats	Reference	Width (bits)
<i>Scalar ALU and Control Formats</i>		
SOP2	page 13-3	32 ¹
SOPK	page 13-6	
SOP1	page 13-8	
SOPC	page 13-11	
SOPP	page 13-13	
<i>Scalar Memory Format</i>		
SMEM	page 13-15	64
<i>Vector ALU Formats</i>		
VOP2	page 13-18	32 ¹
VOP1	page 13-22	32 ¹
VOPC	page 13-26	32 ¹
VOP3 (3 input, one output)	page 13-30	64
VOP3 (3 input, two output)	page 13-36	64
VOP_SDWA	page 13-40	32
VOP_DPP	page 13-42	32
<i>Vector Parameter Interpolation Format</i>		
VINTRP	page 13-44	32
<i>LDS/GDS Format</i>		
DS	page 13-45	64

Table 13.1 Summary of Microcode Formats (Cont.)

Microcode Formats	Reference	Width (bits)
<i>Vector Memory Buffer Formats</i>		
MUBUF MTBUF	page 13-50 page 13-55	64
<i>Vector Memory Image Format</i>		
MIMG	page 13-58	64
<i>Export Formats</i>		
EXP	page 13-62	64
<i>Flat Formats</i>		
FLAT	page 13-63	64

1. This can be 64-bit with a literal constant.

The field-definition tables that accompany the descriptions in the sections below use the following notation.

- *int(2)* — A two-bit field that specifies an unsigned integer value.
- *enum(7)* — A seven-bit field that specifies an enumerated set of values (in this case, a set of up to 2^7 values). The number of valid values can be less than the maximum.

The default value of all fields is zero. Any bitfield not identified is assumed to be reserved.

13.1 Scalar ALU and Control Formats

Scalar Format Two Inputs, One Output

<i>Format</i>	SOP2		
<i>Description</i>	This is a scalar instruction with two inputs and one output. Can be followed by a 32-bit literal constant.		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	SSRC0	[7:0]	enum(8) Source 0. First operand for the instruction. 0 – 101 SGPR0 to SGPR101: Scalar general-purpose registers. 102 FLAT_SCRATCH_LO. 103 FLAT_SCRATCH_HI. 104 XNACK_MASK_LO. Carrizo APU only. 105 XNACK_MASK_HI. Carrizo APU only. 106 VCC_LO: vcc[31:0]. 107 VCC_HI: vcc[63:32]. 108 TBA_LO: Trap handler base address [31:0]. 109 TBA_HI: Trap handler base address [63:32]. 110 TMA_LO: Pointer to data in memory used by trap handler. 111 TMA_HI: Pointer to data in memory used by trap handler. 112 – 123 TIMP0 to TIMP11: Trap handler temporary registers (privileged). 124 M0. Memory register 0. 125 reserved. 126 EXEC_LO: exec[31:0]. 127 EXEC_HI: exec[63:32]. 128 0. 129 – 192 Signed integer 1 to 64. 193 – 208 Signed integer -1 to -16. 209 – 239 reserved. 240 0.5. 241 -0.5. 242 1.0. 243 -1.0. 244 2.0. 245 -2.0. 246 4.0. 247 -4.0. 248 1/(2*PI). 249 – 250 reserved. 251 VCCZ. 252 EXECZ. 253 SCC. 254 reserved. 255 Literal constant.
	SSRC1	[15:8]	enum(8) Source 1. Second operand for instruction. Same codes as for SSRC0, above.
	SDST	[22:16]	enum(7) Scalar destination for instruction. Same codes as for SSRC0, above, except that this can use only codes 0 to 127.

Scalar Format Two Inputs, One Output

OP	[29:23]	enum(7)
----	---------	---------

Opcode.

where the suffix of the instruction specifies the type and size of the result:
D = destination
U = unsigned integer
S = source
SCC = scalar condition code
I = signed integer
B = bitfield

0	S_ADD_U32: D.u = S0.u + S1.u. SCC = carry out.
1	S_SUB_U32: D.u = S0.u - S1.u. SCC = carry out.
2	S_ADD_I32: D.u = S0.i + S1.i. SCC = overflow.
3	S_SUB_I32: D.u = S0.i - S1.i. SCC = overflow.
4	S_ADDC_U32: D.u = S0.u + S1.u + SCC. SCC = carry-out.
5	S_SUBB_U32: D.u = S0.u - S1.u - SCC. SCC = carry-out.
6	S_MIN_I32: D.i = (S0.i < S1.i) ? S0.i : S1.i. SCC = 1 if S0 is min.
7	S_MIN_U32: D.u = (S0.u < S1.u) ? S0.u : S1.u. SCC = 1 if S0 is min.
8	S_MAX_I32: D.i = (S0.i > S1.i) ? S0.i : S1.i. SCC = 1 if S0 is max.
9	S_MAX_U32: D.u = (S0.u > S1.u) ? S0.u : S1.u. SCC = 1 if S0 is max.
10	S_CSELECT_B32: D.u = SCC ? S0.u : S1.u.
11	S_CSELECT_B64: D.u = SCC ? S0.u : S1.u.
12	S_AND_B32: D.u = S0.u & S1.u. SCC = 1 if result is non-zero.
13	S_AND_B64: D.u = S0.u & S1.u. SCC = 1 if result is non-zero.
14	S_OR_B32: D.u = S0.u S1.u. SCC = 1 if result is non-zero.
15	S_OR_B64: D.u = S0.u S1.u. SCC = 1 if result is non-zero.
16	S_XOR_B32: D.u = S0.u ^ S1.u. SCC = 1 if result is non-zero.
17	S_XOR_B64: D.u = S0.u ^ S1.u. SCC = 1 if result is non-zero.
18	S_ANDN2_B32: D.u = S0.u & ~S1.u. SCC = 1 if result is non-zero.
19	S_ANDN2_B64: D.u = S0.u & ~S1.u. SCC = 1 if result is non-zero.
20	S_ORN2_B32: D.u = S0.u ~S1.u. SCC = 1 if result is non-zero.
21	S_ORN2_B64: D.u = S0.u ~S1.u. SCC = 1 if result is non-zero.
22	S_NAND_B32: D.u = ~(S0.u & S1.u). SCC = 1 if result is non-zero.
23	S_NAND_B64: D.u = ~(S0.u & S1.u). SCC = 1 if result is non-zero.
24	S_NOR_B32: D.u = ~(S0.u S1.u). SCC = 1 if result is non-zero.
25	S_NOR_B64: D.u = ~(S0.u S1.u). SCC = 1 if result is non-zero.
26	S_XNOR_B32: D.u = ~(S0.u ^ S1.u). SCC = 1 if result is non-zero.
27	S_XNOR_B64: D.u = ~(S0.u ^ S1.u). SCC = 1 if result is non-zero.
28	S_LSHL_B32: D.u = S0.u << S1.u[4:0]. SCC = 1 if result is non-zero.
29	S_LSHL_B64: D.u = S0.u << S1.u[5:0]. SCC = 1 if result is non-zero.
30	S_LSHR_B32: D.u = S0.u >> S1.u[4:0]. SCC = 1 if result is non-zero.
31	S_LSHR_B64: D.u = S0.u >> S1.u[5:0]. SCC = 1 if result is non-zero.
32	S_ASHR_I32: D.i = signtext(S0.i) >> S1.i[4:0]. SCC = 1 if result is non-zero.
33	S_ASHR_I64: D.i = signtext(S0.i) >> S1.i[5:0]. SCC = 1 if result is non-zero.
34	S_BFM_B32: D.u = ((1 << S0.u[4:0]) - 1) << S1.u[4:0]; bitfield mask.
35	S_BFM_B64: D.u = ((1 << S0.u[5:0]) - 1) << S1.u[5:0]; bitfield mask.
36	S_MUL_I32: D.i = S0.i * S1.i.

Scalar Format Two Inputs, One Output

-
- 37 `S_BFE_U32`: Bit field extract. S0 is data, S1[4:0] is field offset, S1[22:16] is field width. $D.u = (S0.u \gg S1.u[4:0]) \& ((1 \ll S1.u[22:16]) - 1)$. SCC = 1 if result is non-zero.
 - 38 `S_BFE_I32`: Bit field extract. S0 is data, S1[4:0] is field offset, S1[22:16] is field width. $D.i = (S0.u \gg S1.u[4:0]) \& ((1 \ll S1.u[22:16]) - 1)$. SCC = 1 if result is non-zero. Test sign-extended result.
 - 39 `S_BFE_U64`: Bit field extract. S0 is data, S1[4:0] is field offset, S1[22:16] is field width. $D.u = (S0.u \gg S1.u[5:0]) \& ((1 \ll S1.u[22:16]) - 1)$. SCC = 1 if result is non-zero.
 - 40 `S_BFE_I64`: Bit field extract. S0 is data, S1[5:0] is field offset, S1[22:16] is field width. $D.i = (S0.u \gg S1.u[5:0]) \& ((1 \ll S1.u[22:16]) - 1)$. SCC = 1 if result is non-zero. Test sign-extended result.
 - 41 `S_CBRANCH_G_FORK`: Conditional branch using branch stack. Arg0 = compare mask (VCC or any SGPR), Arg1 = 64-bit byte address of target instruction.
 - 42 `S_ABSDIFF_I32`: $D.i = \text{abs}(S0.i \gg S1.i)$. SCC = 1 if result is non-zero.
 - 43 `S_RFE_RESTORE_B64`: Return from exception handler and set: `INST_ATC = S1.U32[0]`

All other values are reserved.

ENCODING	[31:30]	enum(2)
----------	---------	---------

Must be 1 0 .

Scalar Instruction One Inline Constant Input, One Output

<i>Format</i>	SOPK		
<i>Description</i>	This is a scalar instruction with one inline constant input and one output.		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	SIMM16	[15:0]	int(16)
	SDST	[22:16]	enum(7)
	<p>Scalar destination for instruction.</p> <p>Same codes as for SIMM16, above, except that this can use only codes 0 to 127.</p> <p>16-bit integer input for opcode. Signedness is determined by opcode.</p> <p>0 – 101 SGPR0 to SGPR101: Scalar general-purpose registers.</p> <p>102 FLAT_SCRATCH_LO.</p> <p>103 FLAT_SCRATCH_HI.</p> <p>104 XNACK_MASK_LO. Carrizo APU only.</p> <p>105 XNACK_MASK_HI. Carrizo APU only.</p> <p>106 VCC_LO: vcc[31:0].</p> <p>107 VCC_HI: vcc[63:32].</p> <p>108 TBA_LO: Trap handler base address [31:0].</p> <p>109 TBA_HI: Trap handler base address [63:32].</p> <p>110 TMA_LO: Pointer to data in memory used by trap handler.</p> <p>111 TMA_HI: Pointer to data in memory used by trap handler.</p> <p>112 - 123 TTMP0 to TTMP11: Trap handler temporary registers (privileged).</p> <p>124 M0. Memory register 0.</p> <p>125 reserved.</p> <p>126 EXEC_LO: exec[31:0].</p> <p>127 EXEC_HI: exec[63:32].</p> <p>128 0.</p> <p>129 – 192 Signed integer 1 to 64.</p> <p>193 – 208 Signed integer -1 to -16.</p> <p>209 – 239 reserved.</p> <p>240 0.5.</p> <p>241 -0.5.</p> <p>242 1.0.</p> <p>243 -1.0.</p> <p>244 2.0.</p> <p>245 -2.0.</p> <p>246 4.0.</p> <p>247 -4.0.</p> <p>248 1/(2*PI).</p> <p>249 – 250 reserved.</p> <p>251 VCCZ.</p> <p>252 EXECZ.</p> <p>253 SCC.</p> <p>254 reserved.</p> <p>255 Literal constant.</p>		

Scalar Instruction One Inline Constant Input, One Output

OP	[27:23]	enum(5)
<p>Opcode. where the suffix of the instruction specifies the type and size of the result: D = destination U = unsigned integer S = source SCC = scalar condition code I = signed integer B = bitfield</p>		
<p>0 S_MOVK_I32: D.i = signext(SIMM16). 1 S_CMOVK_I32: if (SCC) D.i = signext(SIMM16); else NOP. 2 S_CMPK_EQ_I32: SCC = (D.i == signext(SIMM16)). 3 S_CMPK_LG_I32: SCC = (D.i != signext(SIMM16)). 4 S_CMPK_GT_I32: SCC = (D.i != signext(SIMM16)). 5 S_CMPK_GE_I32: SCC = (D.i >= signext(SIMM16)). 6 S_CMPK_LT_I32: SCC = (D.i < signext(SIMM16)). 7 S_CMPK_LE_I32: SCC = (D.i <= signext(SIMM16)). 8 S_CMPK_EQ_U32: SCC = (D.u == SIMM16). 9 S_CMPK_LG_U32: SCC = (D.u != SIMM16). 10 S_CMPK_GT_U32: SCC = (D.u > SIMM16). 11 S_CMPK_GE_U32: SCC = (D.u >= SIMM16). 12 S_CMPK_LT_U32: SCC = (D.u < SIMM16). 13 S_CMPK_LE_U32: D.u = SCC = (D.u <= SIMM16). 14 S_ADDK_I32: D.i = D.i + signext(SIMM16). SCC = overflow. 15 S_MULK_I32: D.i = D.i * signext(SIMM16). SCC = overflow. 16 S_CBRANCH_I_FORK: Conditional branch using branch-stack. Arg0(sdst) = compare mask (VCC or any SGPR), SIMM16 = signed Dword branch offset relative to next instruction. 17 S_GETREG_B32: D.u = hardware register. Read some or all of a hardware register into the LSBs of D. SIMM16 = {size[4:0], offset[4:0], hwRegId[5:0]}; offset is 0–31, size is 1–32. 18 S_SETREG_B32: hardware register = D.u. Write some or all of the LSBs of D into a hardware register (note that D is a source SGPR). SIMM16 = {size[4:0], offset[4:0], hwRegId[5:0]}; offset is 0–31, size is 1–32. 19 reserved. 20 S_SETREG_IMM32_B32: This instruction uses a 32-bit literal constant. Write some or all of the LSBs of IMM32 into a hardware register. SIMM16 = {size[4:0], offset[4:0], hwRegId[5:0]}; offset is 0–31, size is 1–32. All other values are reserved.</p>		
ENCODING	[31:28]	enum(4)
<p>Must be 1 0 1 1.</p>		

Scalar Instruction One Input, One Output

<i>Format</i>	SOP1		
<i>Description</i>	This is a scalar instruction with one input and one output. Can be followed by a 32-bit literal constant.		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	SSRC0	[7:0]	enum(8)
	Source 0. First operand for the instruction.		
	0 – 101 SGER0 to SGER101: Scalar general-purpose registers.		
	102 FLAT_SCRATCH_LO.		
	103 FLAT_SCRATCH_HI.		
	104 XNACK_MASK_LO. Carrizo APU only.		
	105 XNACK_MASK_HI. Carrizo APU only.		
	106 VCC_LO: vcc[31:0]		
	107 VCC_HI: vcc[63:32]		
	108 TBA_LO: Trap handler base address [31:0].		
	109 TBA_HI: Trap handler base address [63:32].		
	110 TMA_LO: Pointer to data in memory used by trap handler.		
	111 TMA_HI: Pointer to data in memory used by trap handler.		
	112 – 123 TTMP0 to TTMP11: Trap handler temporary registers (privileged).		
	124 M0. Memory register 0.		
	125 reserved.		
	126 EXEC_LO: exec[31:0].		
	127 EXEC_HI: exec[63:32].		
	128 0.		
	129 – 192 Signed integer 1 to 64.		
	193 – 208 Signed integer -1 to -16.		
	209 – 239 reserved.		
	240 0.5.		
	241 -0.5.		
	242 1.0.		
	243 -1.0.		
	244 2.0.		
	245 -2.0.		
	246 4.0.		
	247 -4.0.		
	248 1/(2*PI).		
	249 – 250 reserved.		
	252 EXECZ.		
	253 SCC.		
	254 reserved.		
	255 Literal constant.		

Scalar Instruction One Input, One Output

OP	[15:8]	enum(8)
0	S_MOV_B32:	D.u = S0.u.
1	S_MOV_B64:	D.u = S0.u.
2	S_CMOV_B32:	if(SCC) D.u = S0.u; else NOP.
3	S_CMOV_B64:	if(SCC) D.u = S0.u; else NOP.
4	S_NOT_B32:	D.u = ~S0.u SCC = 1 if result non-zero.
5	S_NOT_B64:	D.u = ~S0.u SCC = 1 if result non-zero.
6	S_WQM_B32:	D.u = WholeQuadMode(S0.u). SCC = 1 if result is non-zero.
7	S_WQM_B64 :	D.u = WholeQuadMode(S0.u). SCC = 1 if result is non-zero.
8	S_BREV_B32:	D.u = S0.u[0:31] (reverse bits).
9	S_BREV_B64:	D.u = S0.u[0:63] (reverse bits).
10	S_BCNT0_I32_B32:	D.i = CountZeroBits(S0.u). SCC = 1 if result is non-zero.
11	S_BCNT0_I32_B64:	D.i = CountZeroBits(S0.u). SCC = 1 if result is non-zero.
12	S_BCNT1_I32_B32:	D.i = CountOneBits(S0.u). SCC = 1 if result is non-zero.
13	S_BCNT1_I32_B64:	D.i = CountOneBits(S0.u). SCC = 1 if result is non-zero.
14	S_FF0_I32_B32:	D.i = FindFirstZero(S0.u) from LSB; if no zeros, return -1.
15	S_FF0_I32_B64:	D.i = FindFirstZero(S0.u) from LSB; if no zeros, return -1.
16	S_FF1_I32_B32:	D.i = FindFirstOne(S0.u) from LSB; if no ones, return -1.
17	S_FF1_I32_B64:	D.i = FindFirstOne(S0.u) from LSB; if no ones, return -1.
18	S_FLBIT_I32_B32:	D.i = FindFirstOne(S0.u) from MSB; if no ones, return -1.
19	S_FLBIT_I32_B64:	D.i = FindFirstOne(S0.u) from MSB; if no ones, return -1.
20	S_FLBIT_I32:	D.i = Find first bit opposite of sign bit from MSB. If S0 == -1, return -1.
21	S_FLBIT_I32_I64:	D.i = Find first bit opposite of sign bit from MSB. If S0 == -1, return -1.
22	S_SEXT_I32_I8:	D.i = signext(S0.i[7:0]).
23	S_SEXT_I32_I16:	D.i = signext(S0.i[15:0]).
24	S_BITSET0_B32:	D.u[S0.u[4:0]] = 0.
25	S_BITSET0_B64:	D.u[S0.u[5:0]] = 0.
26	S_BITSET1_B32:	D.u[S0.u[4:0]] = 1.
27	S_BITSET1_B64:	D.u[S0.u[5:0]] = 1.
28	S_GETPC_B64:	D.u = PC + 4; destination receives the byte address of the next instruction.
29	S_SETPC_B64:	PC = S0.u; S0.u is a byte address of the instruction to jump to.
30	S_SWAPPC_B64:	D.u = PC + 4; PC = S0.u.
31	S_RFE_B64:	Return from Exception; PC = TTMP1,0.
32	S_AND_SAVEEXEC_B64:	D.u = EXEC, EXEC = S0.u & EXEC. SCC = 1 if the new value of EXEC is non-zero.
33	S_OR_SAVEEXEC_B64:	D.u = EXEC, EXEC = S0.u EXEC. SCC = 1 if the new value of EXEC is non-zero.
34	S_XOR_SAVEEXEC_B64:	D.u = EXEC, EXEC = S0.u ^ EXEC. SCC = 1 if the new value of EXEC is non-zero.
35	S_ANDN2_SAVEEXEC_B64:	D.u = EXEC, EXEC = S0.u & ~EXEC. SCC = 1 if the new value of EXEC is non-zero.
36	S_ORN2_SAVEEXEC_B64:	D.u = EXEC, EXEC = S0.u ~EXEC. SCC = 1 if the new value of EXEC is non-zero.
37	S_NAND_SAVEEXEC_B64:	D.u = EXEC, EXEC = ~(S0.u & EXEC). SCC = 1 if the new value of EXEC is non-zero.

Scalar Instruction One Input, One Output

-
- 38 S_NOR_SAVEEXEC_B64: D.u = EXEC, EXEC = ~(S0.u | EXEC). SCC = 1 if the new value of EXEC is non-zero.
 - 39 S_XNOR_SAVEEXEC_B64: D.u = EXEC, EXEC = ~(S0.u ^ EXEC). SCC = 1 if the new value of EXEC is non-zero.
 - 40 S_QUADMASK_B32: D.u = QuadMask(S0.u). D[0] = OR(S0[3:0]), D[1] = OR(S0[7:4]) SCC = 1 if result is non-zero.
 - 41 S_QUADMASK_B64: D.u = QuadMask(S0.u). D[0] = OR(S0[3:0]), D[1] = OR(S0[7:4]) SCC = 1 if result is non-zero. Returns a 64-bit result even though the upper 48 bits are zero.
 - 42 S_MOVRELS_B32: SGPR[D.u] = SGPR[S0.u + M0.u].
 - 43 S_MOVRELS_B64: SGPR[D.u] = SGPR[S0.u + M0.u].
 - 44 S_MOVRELD_B32: SGPR[D.u + M0.u] = SGPR[S0.u].
 - 45 S_MOVRELD_B64: SGPR[D.u + M0.u] = SGPR[S0.u].
 - 46 S_CBRANCH_JOIN: Conditional branch join point. Arg0 = saved CSP value. No dest.
 - 47 reserved.
 - 48 S_ABS_I32: D.i = abs(S0.i). SCC=1 if result is non-zero.
 - 49 S_SET_GPR_IDX_IDX: M0[7:0] = S0.U[7:0]

All other values are reserved.

SDST	[22:16]	enum(7)	Scalar destination for instruction. Same codes as for SSRC0, above, except that this can use only codes 0 to 127.
ENCODING	[31:23]	enum(9)	Must be 1 0 1 1 1 1 1 0 1.

Scalar Instruction Two Inputs, One Comparison

<i>Format</i>	SOPC		
<i>Description</i>	Scalar instruction taking two inputs and producing a comparison result. Can be followed by a 32-bit literal constant.		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	SSRC0	[7:0]	enum(8) Source 0. First operand for the instruction. 0 – 101 SGPR0 to SGPR101: Scalar general-purpose registers. 102 FLAT_SCRATCH_LO. 103 FLAT_SCRATCH_HI. 104 XNACK_MASK_LO. Carrizo APU only. 105 XNACK_MASK_HI. Carrizo APU only. 106 VCC_LO: vcc[31:0] 107 VCC_HI: vcc[63:32] 108 TBA_LO: Trap handler base address [31:0]. 109 TBA_HI: Trap handler base address [63:32]. 110 TMA_LO: Pointer to data in memory used by trap handler. 111 TMA_HI: Pointer to data in memory used by trap handler. 112 – 123 TTMP0 to TTMP11: Trap handler temporary registers (privileged). 124 M0. Memory register 0. 125 reserved. 126 EXEC_LO: exec[31:0]. 127 EXEC_HI: exec[63:32]. 128 0. 129 – 192 Signed integer 1 to 64. 193 – 208 Signed integer -1 to -16. 209 – 239 reserved. 240 0.5. 241 -0.5. 242 1.0. 243 -1.0. 244 2.0. 245 -2.0. 246 4.0. 247 -4.0. 248 1/(2*PI). 249 – 250 reserved. 251 VCCZ. 252 EXECZ. 253 SCC. 254 reserved. 255 Literal constant.
	SSRC1	[15:8]	enum(8) Source 1. Second operand for instruction. Same codes as for SSRC0, above.

Scalar Instruction Two Inputs, One Comparison

OP	[22:16]	enum(8)
	0	S_CMP_EQ_I32: SCC = (S0.i == S1.i).
	1	S_CMP_LG_I32: SCC = (S0.i != S1.i).
	2	S_CMP_GT_I32: SCC = (S0.i > S1.i).
	3	S_CMP_GE_I32: SCC = (S0.i >= S1.i).
	4	S_CMP_LT_I32: SCC = (S0.i < S1.i).
	5	S_CMP_LE_I32: SCC = (S0.i <= S1.i).
	6	S_CMP_EQ_U32: SCC = (S0.u == S1.u).
	7	S_CMP_LG_U32: SCC = (S0.u != S1.u).
	8	S_CMP_GT_U32: SCC = (S0.u > S1.u).
	9	S_CMP_GE_U32: SCC = (S0.u >= S1.u).
	10	S_CMP_LT_U32: SCC = (S0.u < S1.u).
	11	S_CMP_LE_U32: SCC = (S0.u <= S1.u).
	12	S_BITCMP0_B32: SCC = (S0.u[S1.u[4:0]] == 0).
	13	S_BITCMP1_B32: SCC = (S0.u[S1.u[4:0]] == 1).
	14	S_BITCMP0_B64: SCC = (S0.u[S1.u[5:0]] == 0).
	15	S_BITCMP1_B64: SCC = (S0.u[S1.u[5:0]] == 1).
	16	S_SETVSKIP: VSKIP = S0.u[S1.u[4:0]].
	17	S_SET_GPR_IDX_ON: Enable GPR indexing.
	18	S_CMP_EQ_U64: SCC = SCC = (S0.i64 == S1.i64).
	19	S_CMP_NE_U64: SXCCX = (S0 != S1).
ENCODING	[31:23]	enum(9)
	Must be 1 0 1 1 1 1 1 1 0.	

Scalar Instruction One Input, One Special Operation

<i>Format</i>	SOFP		
<i>Description</i>	Scalar instruction taking one inline constant input and performing a special operation (for example: jump).		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	SIMM16	[15:0]	int(16)
16-bit integer input for opcode. Signedness is determined by opcode.			
	OP	[22:16]	enum(8)
	0	S_NOP: do nothing. Repeat NOP 1..8 times based on SIMM16[2:0]. 0 = 1 time, 7 = 8 times.	
	1	S_ENDPGM: end of program; terminate wavefront.	
	2	S_BRANCH: PC = PC + signext(SIMM16 * 4) + 4.	
	3	reserved.	
	4	S_CBRANCH_SCC0: if(SCC == 0) then PC = PC + signext(SIMM16 * 4) + 4; else nop.	
	5	S_CBRANCH_SCC1: if(SCC == 1) then PC = PC + signext(SIMM16 * 4) + 4; else nop.	
	6	S_CBRANCH_VCCZ: if(VCC == 0) then PC = PC + signext(SIMM16 * 4) + 4; else nop.	
	7	S_CBRANCH_VCCNZ: if(VCC != 0) then PC = PC + signext(SIMM16 * 4) + 4; else nop.	
	8	S_CBRANCH_EXECZ: if(EXEC == 0) then PC = PC + signext(SIMM16 * 4) + 4; else nop.	
	9	S_CBRANCH_EXECNZ: if(EXEC != 0) then PC = PC + signext(SIMM16 * 4) + 4; else nop.	
	10	S_BARRIER: Sync waves within a thread group.	
	11	S_SETKILL: Set KILL bit to value of SIMM16[0]. Used primarily for debugging kill wave cmd behavior.	
	12	S_WAITCNT: Wait for count of outstanding lds, vector-memory and export/vmem-write-data to be at or below the specified levels. simm16[3:0] = vmcount, simm16[6:4] = export/mem-write-data count, simm16[12:8] = LGKM_cnt (scalar-mem/GDS/LDS count).	
	13	S_SETHALT: set HALT bit to value of SIMM16[0]. 1=halt, 0=resume. Halt is ignored while priv=1.	
	14	S_SLEEP: Cause a wave to sleep for approximately 64*SIMM16[2:0] clocks.	
	15	S_SETPRIO: User-settable wave priority. 0 = lowest, 3 = highest.	
	16	S_SENDMSG: Send a message.	
	17	S_SENDMSGHALT: Send a message and then HALT.	
	18	S_TRAP: Enter the trap handler. TrapID = SIMM16[7:0]. Wait for all instructions to complete, save {pc_rewind,trapID,pc} into tmp0,1; load TBA into PC, set PRIV=1 and continue.	
	19	S_ICACHE_INV: Invalidate entire L1 I cache.	
	20	S_INCPERFLEVEL: Increment performance counter specified in SIMM16[3:0] by 1.	
	21	S_DECPERFLEVEL: Decrement performance counter specified in SIMM16[3:0] by 1.	
	22	S_TTRACEDATA: Send M0 as user data to thread-trace.	
	23	S_CBRANCH_CDBGSYS: If (conditional_debug_system != 0) then PC = PC + signext(SIMM16 * 4) + 4; else NOP.	

Scalar Instruction One Input, One Special Operation

	24	S_CBRANCH_CDBGUSER: If (conditional_debug_user != 0) then PC = PC + signext(SIMM16 * 4) + 4; else NOP.
	25	S_CBRANCH_CDBGSYS_OR_USER: If (conditional_debug_system conditional_debug_user) then PC = PC + signext(SIMM16 * 4) + 4; else NOP.
	26	S_CBRANCH_CDBGSYS_AND_USER: If (conditional_debug_system && conditional_debug_user) then PC = PC + signext(SIMM16 * 4) + 4; else NOP.
	27	S_ENDPGM_SAVED: End program after context save.
	28	S_SET_GPR_IDX_OFF: Disable GPR index entry.
	29	S_SET_GPR_IDX_MODE: M0[15:12] = SIMM4

ENCODING	[31:23]	enum(9)
	Must be 1 0 1 1 1 1 1 1 1.	

13.2 Scalar Memory Instruction

Scalar Instruction Memory Access

<i>Format</i>	SMEM		
<i>Description</i>	Scalar instruction performing a memory operation on scalar L1 memory. Two Dwords.		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	SBASE	[5:0]	enum(6) Bits [6:1] of an aligned pair of SGPRs specifying {size[15:0], base[47:0]}, where base and size are in Dword units. The low-order bits are in the first SGPR.
	SDATA	[12:6]	enum(7) SGPR that supplies write data, or receives read data. 0 – 101 SGPR0 to SGPR101: Scalar general-purpose registers. 102 FLAT_SCRATCH_LO. 103 FLAT_SCRATCH_HI. 104 XNACK_MASK_LO. Carrizo APU only. 105 XNACK_MASK_HI. Carrizo APU only. 106 VCC_LO: vcc[31:0]. 107 VCC_HI: vcc[63:32]. 108 TBA_LO: Trap handler base address [31:0]. 109 TBA_HI: Trap handler base address [63:32]. 110 TMA_LO: Pointer to data in memory used by trap handler. 111 TMA_HI: Pointer to data in memory used by trap handler. 112 – 123 TTMP0 to TTMP11: Trap handler temporary registers (privileged). All other values reserved.
	reserved	[15:13]	
	GLC	16	enum(1) If set, operation is globally coherent.
	IMM	17	enum(1) Boolean. IMM = 0: Specifies an SGPR address. IMM = 1: Specifies an inline constant offset.

Scalar Instruction Memory Access

OP	[25:18]	enum(8)
0		<code>S_LOAD_DWORD</code> : Read 1 Dword from scalar data cache. If the offset is specified as an SGPR, the SGPR contains an unsigned byte offset (the two LSBs are ignored).
1		<code>S_LOAD_DWORDX2</code> : If the offset is specified as an immediate 20-bit constant, the constant is an unsigned byte offset.
2		<code>S_LOAD_DWORDX4</code> : Read 4 Dwords from scalar data cache. See <code>S_LOAD_DWORD</code> for details on the offset input.
3		<code>S_LOAD_DWORDX8</code> : Read 8 Dwords from scalar data cache. See <code>S_LOAD_DWORD</code> for details on the offset input.
4		<code>S_LOAD_DWORDX16</code> : Read 16 Dwords from scalar data cache. See <code>S_LOAD_DWORD</code> for details on the offset input.
5 - 7		unused.
8		<code>S_BUFFER_LOAD_DWORD</code> : Read one Dword from scalar data cache. See <code>S_LOAD_DWORD</code> for details on the offset input.
9		<code>S_BUFFER_LOAD_DWORDX2</code> : Read two Dwords from scalar data cache. See <code>S_LOAD_DWORD</code> for details on the offset input.
10		<code>S_BUFFER_LOAD_DWORDX4</code> : Read four Dwords from scalar data cache. See <code>S_LOAD_DWORD</code> for details on the offset input.
11		<code>S_BUFFER_LOAD_DWORDX8</code> : Read eight Dwords from scalar data cache. See <code>S_LOAD_DWORD</code> for details on the offset input.
12		<code>S_BUFFER_LOAD_DWORDX16</code> : Read 16 Dwords from scalar data cache. See <code>S_LOAD_DWORD</code> for details on the offset input.
13 - 15		unused.
16		<code>S_STORE_DWORD</code> : Write one Dword to scalar data cache. If the offset is specified as an SGPR, the SGPR contains an unsigned BYTE offset (the two LSBs are ignored). If the offset is specified as an immediate 20-bit constant, the constant is an unsigned byte offset.
17		<code>S_STORE_DWORDX2</code> : Write two Dwords to scalar data cache. See <code>S_STORE_DWORD</code> for details on the offset input.
18		<code>S_STORE_DWORDX4</code> : Write four Dwords to scalar data cache. See <code>S_STORE_DWORD</code> for details on the offset input.
19 - 23		unused.
24		<code>S_BUFFER_STORE_DWORD</code> : Write one Dword to scalar data cache. See <code>S_STORE_DWORD</code> for details on the offset input.
25		<code>S_BUFFER_STORE_DWORDX2</code> : Write two Dwords to scalar data cache. See <code>S_STORE_DWORD</code> for details on the offset input.
26		<code>S_BUFFER_STORE_DWORDX4</code> : Write four Dwords to scalar data cache. See <code>S_STORE_DWORD</code> for details on the offset input.
27 - 31		unused.
32		<code>S_DCACHE_INV</code> : Invalidate the scalar data cache.
33		<code>S_DCACHE_WB</code> : Write back dirty data in the scalar data cache.
34		<code>S_DCACHE_INV_VOL</code> : Invalidate the scalar data cache volatile lines.
35		<code>S_DCACHE_WB_VOL</code> : Write back dirty data in the scalar data cache volatile lines.
36		<code>S_MEMTIME</code> : Return current 64-bit timestamp.
37		<code>S_MEMREALTIME</code> : Return current 64-bit RTC.
38		<code>S_ATC_PROBE</code> : Probe or prefetch an address into the SQC data cache.
39		<code>S_ATC_PROBE_BUFFER</code> : Probe or prefetch an address into the SQC data cache.
		All other values are reserved.

Scalar Instruction Memory Access

ENCODING	[31:26]	enum(6)
	Must be 1 1 0 0 0 0.	
OFFSET	[51:32]	enum(6)
	IMM = 0: Specifies an SGPR address that supplies a byte offset for the memory operation (see enumeration).	
	IMM = 1: Specifies a 20-bit unsigned byte offset.	
reserved	[63:52]	

13.3 Vector ALU instructions

Vector Instruction Two Inputs, One Output

<i>Format</i>	VOP2		
<i>Description</i>	Vector instruction taking two inputs and producing one output. Can be followed by a 32-bit literal constant.		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	SRC0	[8:0]	enum(9) First operand for instruction. Source 0. First operand for the instruction. 0 – 101 SGPR0 to SGPR101: Scalar general-purpose registers. 102 FLAT_SCRATCH_LO. 103 FLAT_SCRATCH_HI. 104 XNACK_MASK_LO. Carrizo APU only. 105 XNACK_MASK_HI. Carrizo APU only. 106 VCC_LO: vcc[31:0]. 107 VCC_HI: vcc[63:32]. 108 TBA_LO: Trap handler base address [31:0]. 109 TBA_HI: Trap handler base address [63:32]. 110 TMA_LO: Pointer to data in memory used by trap handler. 111 TMA_HI: Pointer to data in memory used by trap handler. 112 – 123 TTMP0 to TTMP11: Trap handler temporary registers (privileged). 124 M0. Memory register 0. 125 reserved. 126 EXEC_LO: exec[31:0]. 127 EXEC_HI: exec[63:32]. 128 0. 129 – 192 Signed integer 1 to 64. 193 – 208 Signed integer -1 to -16. 209 – 239 reserved. 240 0.5 241 -0.5. 242 1.0. 243 -1.0. 244 2.0. 245 -2.0. 246 4.0. 247 -4.0. 248 1/(2*PI). 249 – 250 reserved. 251 VCCZ. 252 EXECZ. 253 SCC. 254 LDS direct. 255 Literal constant. 256 – 511 Vector General-Purpose Registers (VGPRs) 0 – 255.
	VSRC1	[16:9]	enum(8) Second operand for instruction. 0 - 255 Vector General-Purpose Registers (VGPRs) 0 - 255.

Vector Instruction Two Inputs, One Output

VDST	[24:17]	enum(8)
	Destination for instruction. 0 - 255 Vector General-Purpose Registers (VGPRs) 0 - 255.	
OP	[30:25]	enum(7)
	0	V_CNDMASK_B32: D.u = VCC[i] ? S1.u : S0.u (i = threadID in wave); VOP3: specify VCC as a scalar GPR in S2.
	1	V_ADD_F32: D.f = S0.f + S1.f.
	2	V_SUB_F32: D.f = S0.f - S1.f.
	3	V_SUBREV_F32: D.f = S1.f - S0.f. SQ translates to V_ADD_F32.
	4	V_MUL_LEGACY_F32: D.f = S0.f * S1.f (DX9 rules, 0.0*x = 0.0).
	5	V_MUL_F32: D.f = S0.f * S1.f.
	6	V_MUL_I32_I24: D.i = S0.i[23:0] * S1.i[23:0].
	7	V_MUL_HI_I32_I24: D.i = (S0.i[23:0] * S1.i[23:0])>>32.
	8	V_MUL_U32_U24: D.u = S0.u[23:0] * S1.u[23:0].
	9	V_MUL_HI_U32_U24: D.i = (S0.u[23:0] * S1.u[23:0])>>32.
	10	V_MIN_F32: D.f = (S0.f < S1.f ? S0.f : S1.f).
	11	V_MAX_F32: D.f = (S0.f >= S1.f ? S0.f : S1.f).
	12	V_MIN_I32: D.i = min(S0.i, S1.i).
	13	V_MAX_I32: D.i = max(S0.i, S1.i).
	14	V_MIN_U32: D.u = min(S0.u, S1.u).
	15	V_MAX_U32: D.u = max(S0.u, S1.u).
	16	V_LSHRREV_B32: D.u = S1.u >> S0.u[4:0]. The vacated bits are set to zero. SQ translates this to an internal SP opcode.
	17	V_ASHRREV_I32: D.i = signtext(S1.i) >> S0.i[4:0]. The vacated bits are set to the sign bit of the input value. SQ translates this to an internal SP opcode.
	18	V_LSHLREV_B32: D.u = S1.u << S0.u[4:0]. SQ translates this to an internal SP opcode.
	19	V_AND_B32: D.u = S0.u & S1.u. Input and output modifiers not supported.
	20	V_OR_B32: D.u = S0.u S1.u. Input and output modifiers not supported.
	21	V_XOR_B32: D.u = S0.u ^ S1.u. Input and output modifiers not supported.
	22	V_MAC_F32: D.f = S0.f * S1.f + D.f. SQ translates to V_MAD_F32.
	23	V_MADMK_F32: D.f = S0.f * K + S1.f; K is a 32-bit inline constant. This opcode cannot use the VOP3 encoding and cannot use input/output modifiers. SQ translates to V_MAD_F32.
	24	V_MADAK_F32: D.f = S0.f * S1.f + K; K is a 32-bit inline constant. This opcode cannot use the VOP3 encoding and cannot use input/output modifiers. SQ translates to V_MAD_F32.
	25	V_ADD_U32: D.u = S0.u + S1.u; \nVCC[threadId] = (S0.u + S1.u >= 0x80000000ULL ? 1 : 0) is an unsigned overflow or carry-out for V_ADDC_U32. In VOP3, the VCC destination can be an arbitrary SGPR-pair.
	26	V_SUB_U32: D.u = S0.u - S1.u; VCC[threadId] = (S1.u > S0.u ? 1 : 0) is an unsigned overflow or carry-out for V_SUBB_U32. In VOP3, the VCC destination can be an arbitrary SGPR-pair.
	27	V_SUBREV_U32: D.u = S1.u - S0.u; \nVCC[threadId] = (S0.u > S1.u ? 1 : 0) is an unsigned overflow or carry-out for V_SUBB_U32. In VOP3, the VCC destination can be an arbitrary SGPR-pair. SQ translates this to V_SUB_U32 with reversed operands.
	28	V_ADDC_U32: D.u = S0.u + S1.u + VCC[threadId]; \nVCC[threadId] = (S0.u + S1.u + VCC[threadId] >= 0x80000000ULL ? 1 : 0) is an unsigned overflow. In VOP3, the VCC destination can be an arbitrary SGPR-pair, and the VCC source comes from the SGPR-pair at S2.u.

Vector Instruction Two Inputs, One Output

- 29 V_SUBB_U32 : $D.u = S0.u - S1.u - VCC[threadld]$; $\backslash n VCC[threadld] = (S1.u + VCC[threadld] > S0.u ? 1 : 0)$ is an unsigned overflow. In VOP3, the VCC destination can be an arbitrary SGPR-pair, and the VCC source comes from the SGPR-pair at $S2.u$.
- 30 $V_SUBBREV_U32$: $D.u = S1.u - S0.u - VCC[threadld]$; $\backslash n VCC[threadld] = (S1.u + VCC[threadld] > S0.u ? 1 : 0)$ is an unsigned overflow. In VOP3, the VCC destination can be an arbitrary SGPR-pair, and the VCC source comes from the SGPR-pair at $S2.u$. SQ translates to V_SUBB_U32 . SQ translates this to V_SUBREV_U32 with reversed operands.
- 31 V_ADD_F16 : $D.f16 = S0.f16 + S1.f16$. Supports denormals, round mode, exception flags, saturation.
- 32 V_SUB_F16 : $D.f16 = S0.f16 - S1.f16$. Supports denormals, round mode, exception flags, saturation. SQ translates to V_ADD_F16 .
- 33 V_SUBREV_F16 : $D.f16 = S1.f16 - S0.f16$. Supports denormals, round mode, exception flags, saturation. SQ translates to V_ADD_F16 .
- 34 V_MUL_F16 : $D.f16 = S0.f16 * S1.f16$. Supports denormals, round mode, exception flags, saturation.
- 35 V_MAC_F16 : $D.f16 = S0.f16 * S1.f16 + D.f16$. Supports round mode, exception flags, saturation. SQ translates this to V_MAD_F16 .
- 36 V_MADMK_F16 : $D.f16 = S0.f16 * K.f16 + S1.f16$; K is a 16-bit inline constant stored in the following literal Dword. This opcode cannot use the VOP3 encoding and cannot use input/output modifiers. Supports round mode, exception flags, saturation. SQ translates this to V_MAD_F16 .
- 37 V_MADAK_F16 : $D.f16 = S0.f16 * S1.f16 + K.f16$; K is a 16-bit inline constant stored in the following literal Dword. This opcode cannot use the VOP3 encoding and cannot use input/output modifiers. Supports round mode, exception flags, saturation. SQ translates this to V_MAD_F16 .
- 38 V_ADD_U16 : $D.u16 = S0.u16 + S1.u16$. Supports saturation (unsigned 16-bit integer domain).
- 39 V_SUB_U16 : $D.u16 = S0.u16 - S1.u16$. Supports saturation (unsigned 16-bit integer domain).
- 40 V_SUBREV_U16 : $D.u16 = S1.u16 - S0.u16$. Supports saturation (unsigned 16-bit integer domain). SQ translates this to V_SUB_U16 with reversed operands.
- 41 $V_MUL_LO_U16$: $D.u16 = S0.u16 * S1.u16$. Supports saturation (unsigned 16-bit integer domain).
- 42 $V_LSHLREV_B16$: $D.u[15:0] = S1.u[15:0] \ll S0.u[3:0]$. SQ translates this to an internal SP opcode.
- 43 $V_LSHRREV_B16$: $D.u[15:0] = S1.u[15:0] \gg S0.u[3:0]$. The vacated bits are set to zero. SQ translates this to an internal SP opcode.
- 44 V_ASHREV_I16 : $D.i[15:0] = \text{signext}(S1.i[15:0]) \gg S0.i[3:0]$. The vacated bits are set to the sign bit of the input value. SQ translates this to an internal SP opcode.
- 45 V_MAX_F16 : $D.f16 = \max(S0.f16, S1.f16)$. IEEE compliant. Supports denormals, round mode, exception flags, saturation.
- 46 V_MIN_F16 : $D.f16 = \min(S0.f16, S1.f16)$. IEEE compliant. Supports denormals, round mode, exception flags, saturation.
- 47 V_MAX_U16 : $D.u[15:0] = \max(S0.u[15:0], S1.u[15:0])$.
- 48 V_MAX_I16 : $D.i[15:0] = \max(S0.i[15:0], S1.i[15:0])$.
- 49 V_MIN_U16 : $D.u[15:0] = \min(S0.u[15:0], S1.u[15:0])$.
- 50 V_MIN_I16 : $D.i[15:0] = \min(S0.i[15:0], S1.i[15:0])$.
- 51 V_LDEXP_F16 : $D.f16 = S0.f16 * (2 ** S1.i16)$.
- All other values are reserved.
-

Vector Instruction Two Inputs, One Output

Encode	31	enum(1)
	Must be 0.	

Vector Instruction One Input, One Output

<i>Format</i>	VO_{PI}		
<i>Description</i>	Vector instruction taking one input and producing one output. Can be followed by a 32-bit literal constant.		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	SRC0	[8:0]	enum(9)
			First operand for instruction.
			Source 0. First operand for the instruction.
			0 – 101 SGPR0 to SGPR101: Scalar general-purpose registers.
			102 FLAT_SCRATCH_LO.
			103 FLAT_SCRATCH_HI.
			104 XNACK_MASK_LO. Carrizo APU only.
			105 XNACK_MASK_HI. Carrizo APU only.
			106 VCC_LO: vcc[31:0].
			107 VCC_HI: vcc[63:32].
			108 TBA_LO: Trap handler base address [31:0].
			109 TBA_HI: Trap handler base address [63:32].
			110 TMA_LO: Pointer to data in memory used by trap handler.
			111 TMA_HI: Pointer to data in memory used by trap handler.
			112 – 123 TTMP0 to TTMP11: Trap handler temporary registers (privileged).
			124 M0. Memory register 0.
			125 reserved.
			126 EXEC_LO: exec[31:0].
			127 EXEC_HI: exec[63:32].
			128 0.
			129 – 192 Signed integer 1 to 64.
			193 – 208 Signed integer -1 to -16.
			209 – 239 reserved.
			240 0.5.
			241 -0.5.
			242 1.0.
			243 -1.0.
			244 2.0.
			245 -2.0.
			246 4.0.
			247 -4.0.
			248 1/(2*PI).
			249 – 250 reserved.
			251 VCCZ.
			252 EXECZ.
			253 SCC.
			254 LDS direct.
			255 Literal constant.
			256 – 511 Vector General-Purpose Registers (VGPRs) 0 – 255.

Vector Instruction One Input, One Output (Cont.)

OP	[16:9]	enum(8)
	0	V_NOP: do nothing.
	1	V_MOV_B32: D.u = S0.u. Input and output modifiers not supported; this is an untyped operation.
	2	V_READFIRSTLANE_B32: Copy one VGPR value to one SGPR. Dst = SGPR-dest, Src0 = source data (VGPR# or M0(lds-direct)), Lane# = FindFirst1fromLSB(exec) (lane = 0 if exec is zero). Ignores exec mask for the access. SQ translates to V_READLANE_B32. Input and output modifiers not supported; this is an untyped operation
	3	V_CVT_I32_F64: D.i = (int)S0.d. Out-of-range floating point values (including infinity) saturate. NaN is converted to 0.
	4	V_CVT_F64_I32: D.d = (double)S0.i.
	5	V_CVT_F32_I32: D.f = (float)S0.i.
	6	V_CVT_F32_U32: D.f = (float)S0.u.
	7	V_CVT_U32_F32: D.u = (unsigned)S0.f. Out-of-range floating point values (including infinity) saturate. NaN is converted to 0.
	8	V_CVT_I32_F32: D.i = (int)S0.f. Out-of-range floating point values (including infinity) saturate. NaN is converted to 0.
	9	reserved.
	10	V_CVT_F16_F32: D.f16 = flt32_to_flt16(S0.f). Supports input modifiers and creates FP16 denormals when appropriate.
	11	V_CVT_F32_F16: D.f = flt16_to_flt32(S0.f16). FP16 denormal inputs are always accepted.
	12	V_CVT_RPI_I32_F32: D.i = (int)floor(S0.f + 0.5).
	13	V_CVT_FLR_I32_F32: D.i = (int)floor(S0.f).
	14	V_CVT_OFF_F32_I4: 4-bit signed int to 32-bit float. For interpolation in shader.
	15	V_CVT_F32_F64: D.f = (float)S0.d.
	16	V_CVT_F64_F32: D.d = (double)S0.f.
	17	V_CVT_F32_UBYTE0: D.f = (float)(S0.u[7:0]).
	18	V_CVT_F32_UBYTE1: D.f = (float)(S0.u[15:8]).
	19	V_CVT_F32_UBYTE2: D.f = (float)(S0.u[23:16]).
	20	V_CVT_F32_UBYTE3: D.f = (float)(S0.u[31:24]).
	21	V_CVT_U32_F64: D.u = (uint)S0.d. Out-of-range floating point values (including infinity) saturate. NaN is converted to 0.
	22	V_CVT_F64_U32: D.d = (double)S0.u.
	23	V_TRUNC_F64: D.d = trunc(S0.d), return integer part of S0.d.
	24	V_CEIL_F64: D.d = trunc(S0.d); if (S0.d > 0.0 && S0.d != D.d), D.d += 1.0.
	25	V_RNDNE_F64: D.d = round_nearest_even(S0.d).
	26	V_FLOOR_F64: D.d = trunc(S0.d); if (S0.d < 0.0 && S0.d != D.d), D.d += -1.0.
	27	V_FRACT_F32: D.f = S0.f - floor(S0.f).
	28	V_TRUNC_F32: D.f = trunc(S0.f), return integer part of S0.
	29	V_CEIL_F32: D.f = trunc(S0.f). If (S0 > 0.0 && S0 != D), D += 1.0.
	30	V_RNDNE_F32: D.f = round_nearest_even(S0.f).
	31	V_FLOOR_F32: D.f = trunc(S0.f); if ((S0.f < 0.0) && (S0.f != D.f)), then D.f += -1.0.
	32	V_EXP_F32: D.f = pow(2.0, S0.f).
	33	V_LOG_F32: D.f = log2(S0.f). Base 2 logarithm.
	34	V_RCP_F32: D.f = 1.0 / S0.f. Reciprocal with IEEE rules and < 1ulp error.
	35	V_RCP_IFLAG_F32: D.f = 1.0 / S0.f. Reciprocal intended for integer division, can raise integer DIV_BY_ZERO exception but cannot raise floating-point exceptions.
	36	V_RSQ_F32: D.f = 1.0 / sqrt(S0.f). Reciprocal square root with IEEE rules.

Vector Instruction One Input, One Output (Cont.)

-
- 37 V_RCP_F64: $D.d = 1.0 / (S0.d)$.
- 38 V_RSQ_F64: $D.f = 1.0 / \text{sqrt}(S0.f)$. Reciprocal square root with IEEE rules.
- 39 V_SQRT_F32: $D.f = \text{sqrt}(S0.f)$.
- 40 V_SQRT_F64: $D.d = \text{sqrt}(S0.d)$.
- 41 V_SIN_F32: $D.f = \sin(S0.f * 2 * \text{PI})$. Valid range of S0.f is [-256.0, +256.0]. Out of range input results in float 0.0.
- 42 V_COS_F32: $D.f = \cos(S0.f * 2 * \text{PI})$. Valid range of S0.f is [-256.0, +256.0]. Out of range input results in float 0.0.
- 43 V_NOT_B32: $D.u = \sim S0.u$. Input and output modifiers not supported.
- 44 V_BFREV_B32: $D.u[31:0] = S0.u[0:31]$, bitfield reverse. Input and output modifiers not supported.
- 45 V_FFBL_U32: $D.u = \text{position of first 1 in } S0.u \text{ from MSB}$; $D.u=0xFFFFFFFF$ if $S0==0$.
- 46 V_FFBL_B32: $D.u = \text{position of first 1 in } S0.u \text{ from LSB}$; $D.u=0xFFFFFFFF$ if $S0==0$.
- 47 V_FFBH_I32: $D.u = \text{position of first bit different from sign bit in } S0.i \text{ from MSB}$; $D.u=0xFFFFFFFF$ if $S0==0$ or $0.i == 0xFFFFFFFF$.
- 48 V_FREXP_EXP_I32_F64: See V_FREXP_EXP_I32_F32.
- 49 V_FREXP_MANT_F64: See V_FREXP_MANT_F32.
- 50 V_FRACT_F64: See V_FRACT_F32.
- 51 V_FREXP_EXP_I32_F32: If $(S0.f == \text{INF} || S0.f == \text{NAN})$, then $D.i = 0$; else $D.i = \text{TwosComplement}(\text{Exponent}(S0.f) - 127 + 1)$. Returns exponent of single precision float input, such that $S0.f = \text{significand} * (2 ** \text{exponent})$. See also FREXP_MANT_F32, which returns the significand.
- 52 V_FREXP_MANT_F32: if $(S0.f == \text{INF} || S0.f == \text{NAN})$ then $D.f = S0.f$; else $D.f = \text{Mantissa}(S0.f)$. Result range is in $(-1.0, -0.5][0.5, 1.0)$ in normal cases. Returns binary significand of single precision float input, such that $S0.f = \text{significand} * (2 ** \text{exponent})$. See also FREXP_EXP_I32_F32, which returns integer exponent.
- 53 V_CLREXCP: Clear wave's exception state in SIMD(SP).
- 54 V_MOVRELD_B32: $\text{VGPR}[D.u + M0.u] = \text{VGPR}[S0.u]$. Input and output modifiers not supported; this is an untyped operation.
- 55 V_MOVRELS_B32: $\text{VGPR}[D.u] = \text{VGPR}[S0.u + M0.u]$. Input and output modifiers not supported; this is an untyped operation.
- 56 V_MOVRELS_D_B32: $\text{VGPR}[D.u + M0.u] = \text{VGPR}[S0.u + M0.u]$. Input and output modifiers not supported; this is an untyped operation. SQ translates to V_MOV_B32.
- 57 V_CVT_F16_U16: $D.f16 = \text{uint16_to_flt16}(S.u16)$. Supports denormals, rounding, exception flags and saturation.
- 58 V_CVT_F16_I16: $D.f16 = \text{int16_to_flt16}(S.i16)$. Supports denormals, rounding, exception flags and saturation.
- 59 V_CVT_U16_F16: $D.u16 = \text{flt16_to_uint16}(S.f16)$. Supports rounding, exception flags and saturation.
- 60 V_CVT_I16_F16: $D.i16 = \text{flt16_to_int16}(S.f16)$. Supports rounding, exception flags and saturation.
- 61 V_RCP_F16: if $(S0.f16 == 1.0f)$, $D.f16 = 1.0f$; else $D.f16 = \text{ApproximateRecip}(S0.f16)$.
- 62 V_SQRT_F16: if $(S0.f16 == 1.0f)$ $D.f16 = 1.0f$; else $D.f16 = \text{ApproximateSqrt}(S0.f16)$.
- 63 V_RSQ_F16: if $(S0.f16 == 1.0f)$ $D.f16 = 1.0f$; else $D.f16 = \text{ApproximateRecipSqrt}(S0.f16)$.
- 64 V_LOG_F16: if $(S0.f16 == 1.0f)$ $D.f16 = 0.0f$; else $D.f16 = \text{ApproximateLog2}(S0.f16)$.
-

Vector Instruction One Input, One Output (Cont.)

-
- 65 V_EXP_F16: if(S0.f16 == 0.0f) D.f16 = 1.0f; else D.f16 = Approximate2ToX(S0.f16).
 - 66 V_FREXP_MANT_F16: if(S0.f16 == +-INF || S0.f16 == NAN) D.f16 = S0.f16; else D.f16 = mantissa(S0.f16). Result range is (-1.0,-0.5][0.5,1.0). C math library frexp function. Returns binary significand of half precision float input, such that the original single float = significand * (2 ** exponent).
 - 67 V_FREXP_EXP_I16_F16: if(S0.f16 == +-INF || S0.f16 == NAN) D.i16 = 0; else D.i16 = 2s_complement(exponent(S0.f16) - 15 + 1). C math library frexp function. Returns exponent of half precision float input, such that the original single float = significand * (2 ** exponent).
 - 68 V_FLOOR_F16: D.f16 = trunc(S0.f16); if(S0.f16 < 0.0f && S0.f16 != D.f16) then D.f16 -= 1.0f.
 - 69 V_CEIL_F16: D.f16 = trunc(S0.f16); if(S0.f16 > 0.0f && S0.f16 != D.f16) then D.f16 += 1.0f.
 - 70 V_TRUNC_F16: D.f16 = trunc(S0.f16).\n\nRound-to-zero semantics.
 - 71 V_RNDNE_F16: D.f16 = FLOOR(S0.f16 + 0.5f); if(floor(S0.f16) is even && fract(S0.f16) == 0.5f) then D.f16 -= 1.0f. Round-to-nearest-even semantics.
 - 72 V_FRACT_F16: D.f16 = S0.f16 + -floor(S0.f16).
 - 73 V_SIN_F16: D.f16 = sin(S0.f16 * 2 * PI).
 - 74 V_COS_F16: D.f16 = cos(S0.f16 * 2 * PI).
 - 75 V_EXP_LEGACY_F32: D.f = pow(2.0, S0.f) with legacy semantics.
 - 76 V_LOG_LEGACY_F32: D.f = log2(S0.f). Base 2 logarithm with legacy semantics.
- All other values reserved.

VDST	[24:17	enum(8)	Destination for instruction. 0 – 255Vector General-Purpose Registers (VGPRs) 0 – 255.
ENCODE	[31:25]	enum(7)	Must be 0 1 1 1 1 1 1.

Vector Instruction Two Inputs, One Comparison Result

Format **VOPC**

Description Vector instruction taking two inputs and producing a comparison result. Can be followed by a 32-bit literal constant.

Vector Comparison operations are divided into three groups:

- those which can use any one of 16 comparison operations,
- those which can use any one of 8, and
- those which have only a single comparison operation.

The final opcode number is determined by adding the base for the opcode family plus the offset from the compare op.

Every compare instruction writes a result to VCC (for VOPC) or an SGPR (for VOP3).

Additionally, every compare instruction has a variant that also writes to the EXEC mask.

The destination of the compare result is always VCC when encoded using the VOPC format, and can be an arbitrary SGPR when encoded in the VOP3 format.

Vector Instruction Two Inputs, One Comparison Result

<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	SRC0	[8:0]	enum(9)
			First operand for instruction. Source 0. First operand for the instruction. 0 – 101 SGPR0 to SGPR101: Scalar general-purpose registers. 102 FLAT_SCRATCH_LO. 103 FLAT_SCRATCH_HI. 104 XNACK_MASK_LO. Carrizo APU only. 105 XNACK_MASK_HI. Carrizo APU only. 106 VCC_LO: vcc[31:0]. 107 VCC_HI: vcc[63:32]. 108 TBA_LO: Trap handler base address [31:0]. 109 TBA_HI: Trap handler base address [63:32]. 110 TMA_LO: Pointer to data in memory used by trap handler. 111 TMA_HI: Pointer to data in memory used by trap handler. 112 – 123 TTMP0 to TTMP11: Trap handler temporary registers (privileged). 124 M0. Memory register 0. 125 reserved. 126 EXEC_LO: exec[31:0]. 127 EXEC_HI: exec[63:32]. 128 0. 129 – 192: Signed integer 1 to 64. 193 – 208: Signed integer -1 to -16. 209 – 239: reserved. 240 0.5. 241 -0.5. 242 1.0. 243 -1.0. 244 2.0. 245 -2.0. 246 4.0. 247 -4.0. 248 1/(2*PI). 249 – 250 reserved. 251 VCCZ. 252 EXECZ. 253 SCC. 254 LDS direct. 255 Literal constant. 256 - 511 Vector General-Purpose Registers (VGPRs) 0 - 255.
	VSRC1	[16:9]	enum(8)
			Second operand for instruction. 0 – 255 Vector General-Purpose Registers (VGPRs) 0 - 255.

Vector Instruction Two Inputs, One Comparison Result

OP [24:17] enum(8)

Sixteen Compare Operations (OP16)Compare Opcode

<u>Operation</u>	<u>Offset</u>	<u>Description</u>
F	0	D.u = 0
LT	1	D.u = (S0 < S1)
EQ	2	D.u = (S0 == S1)
LE	3	D.u = (S0 <= S1)
GT	4	D.u = (S0 > S1)
LG	5	D.u = (S0 <> S1)
GE	6	D.u = (S0 >= S1)
O	7	D.u = (!isNaN(S0) && !isNaN(S1))
U	8	D.u = (!isNaN(S0) !isNaN(S1))
NGE	9	D.u = !(S0 >= S1)
NLG	10	D.u = !(S0 <> S1)
NGT	11	D.u = !(S0 > S1)
NLE	12	D.u = !(S0 <= S1)
NEQ	13	D.u = !(S0 == S1)
NLT	14	D.u = !(S0 < S1)
TRU	15	D.u = 1

Eight Compare Operations (OP8)Compare Opcode

<u>Operation</u>	<u>Offset</u>	<u>Description</u>
F	0	D.u = 0
LT	1	D.u = (S0 < S1)
EQ	2	D.u = (S0 == S1)
LE	3	D.u = (S0 <= S1)
GT	4	D.u = (S0 > S1)
LG	5	D.u = (S0 <> S1)
GE	6	D.u = (S0 >= S1)
TRU	7	D.u = 1

Vector Instruction Two Inputs, One Comparison Result

Single Vector Compare Operations

		<u>Opcode</u>	
<u>Opcode Family</u>	<u>Base</u>	<u>Description</u>	
Reserved	0x00		
CMP_CLASS_F32	0x10	none	
CMPX_CLASS_F32	0x11	none	
CMP_CLASS_F64	0x12	none	
CMPX_CLASS_F64	0x13	none	
CMP_CLASS_F16	0x14	none	
CMPX_CLASS_F16	0x15	none	
Reserved	0x16-0x1F		
CMP_F16	0x20	OP16	
CMPX_F16	0x30	OP16	
CMP_F32	0x40	OP16	
CMPX_F32	0x50	OP16	
CMP_F64	0x60	OP16	
CMPX_F64	0x70	OP16	
Reserved	0x80-0x98		
CMP_I16	0xA0	OP8	
CMP_U16	0xA8	OP8	
CMPX_I16	0xB0	OP8	
CMPX_U16	0xB8	OP8	
CMP_I32	0xC0	OP8	
CMP_U32	0xC8	OP8	
CMPX_I32	0xD0	OP8	
CMPX_U32	0xD8	OP8	
CMP_I64	0xE0	OP8	
CMP_U64	0xE8	OP8	
CMPX_I64	0xF0	OP8	
CMPX_U64	0xF8	OP8	
ENCODE	[31:25]	enum(7)	
Must be 0 1 1 1 1 1 0.			

Vector Instruction, Three Inputs, One Output (VOP3a)

Format	VOP3a		
Description	Vector instruction taking three inputs and producing one output.		
Opcodes	Field Name	Bits	Format
	VDST	[7:0]	enum(8) Destination for instruction in the Vector General-Purpose Registers (VGPR [255:0]). For V_CMP instructions, this field specifies the SGPR or VCC that receives the result of the comparison.
	ABS	[10:8]	enum(3) If ABS[N] is set, take the floating-point absolute value of the N'th input operand. This is applied before negation.
	reserved	[14:11]	Reserved.
	CLAMP	15	enum(1) If set, clamp output to [0.0, 1.0]. Applied after output modifier.
	OP	[25:16]	enum(10) 0 – 255 Are the VOPC opcodes when VOP3 encoding is required. For example, 0 + V_CMP_F_F32 generates the VOP3 version of V_CMP_F_F32. 256 – 319 Are the VOP2 opcodes with 256 added to their respective numbers. 320 – 447 Are the VOP1 opcodes with 320 added to their respective numbers. 448 V_MAD_LEGACY_F32: D.f = S0.f * S1.f + S2.f (DX9 rules, 0.0*x = 0.0). 449 V_MAD_F32: D.f = S0.f * S1.f + S2.f. 450 V_MAD_I32_I24: D.i = S0.i[23:0] * S1.i[23:0] + S2.iD.i 451 V_MAD_U32_U24: D.u = S0.u[23:0] * S1.u[23:0] + S2.u. 452 V_CUBEID_F32: D.f = cubemap face ID ({0.0, 1.0, ..., 5.0}). XYZ coordinate is given in (S0.f, S1.f, S2.f). 453 V_CUBESC_F32: D.f = cubemap S coordinate. XYZ coordinate is given in (S0.f, S1.f, S2.f). 454 V_CUBETC_F32: D.f = cubemap T coordinate. XYZ coordinate is given in (S0.f, S1.f, S2.f). 455 V_CUBEMA_F32: D.f = 2.0 * cubemap major axis. XYZ coordinate is given in (S0.f, S1.f, S2.f). 456 V_BFE_U32: D.u = (S0.u >> S1.u[4:0]) & ((1 << S2.u[4:0]) - 1). Bitfield extract with S0 = data, S1 = field_offset, S2 = field_width. 457 V_BFE_I32: D.i = (S0.i >> S1.u[4:0]) & ((1 << S2.u[4:0]) - 1). Bitfield extract with S0 = data, S1 = field_offset, S2 = field_width 458 V_BFI_B32: D.u = (S0.u & S1.u) (~S0.u & S2.u); bitfield insert. 459 V_FMA_F32: D.f = S0.f * S1.f + S2.f. 460 V_FMA_F64: D.d = S0.d * S1.d + S2.d. 461 V_LERP_U8: D.u = ((S0.u[31:24] + S1.u[31:24] + S2.u[24]) >> 1) << 24\nD.u += ((S0.u[23:16] + S1.u[23:16] + S2.u[16]) >> 1) << 16;\nD.u += ((S0.u[15:8] + S1.u[15:8] + S2.u[8]) >> 1) << 8;\nD.u += ((S0.u[7:0] + S1.u[7:0] + S2.u[0]) >> 1). Unsigned 8-bit pixel average on packed unsigned bytes (linear interpolation). S2 acts as a round mode; if set, 0.5 rounds up, otherwise 0.5 truncates.

Vector Instruction, Three Inputs, One Output (VOP3a)

-
- 462 V_ALIGNBIT_B32: D.u = $\{(S0,S1) \gg S2.u[4:0]\} \& 0xFFFFFFFF$.
- 463 V_ALIGNBYTE_B32: D.u = $\{(S0,S1) \gg (8*S2.u[4:0])\} \& 0xFFFFFFFF$.
- 464 V_MIN3_F32: D.f = min(S0.f, S1.f, S2.f).
- 465 V_MIN3_I32: D.i = min(S0.i, S1.i, S2.i).
- 466 V_MIN3_U32: 0x153 D.u = min(S0.u, S1.u, S2.u).
- 467 V_MAX3_F32: D.f = max(S0.f, S1.f, S2.f).
- 468 V_MAX3_I32: D.i = max(S0.i, S1.i, S2.i).
- 469 V_MAX3_U32: D.u = max(S0.u, S1.u, S2.u).
- 470 V_MED3_F32: D.f = median(S0.f, S1.f, S2.f).
- 471 V_MED3_I32: D.i = median(S0.i, S1.i, S2.i).
- 472 V_MED3_U32: D.u = median(S0.u, S1.u, S2.u).
- 473 V_SAD_U8: D.u = $\text{abs}(S0.i[31:24] - S1.i[31:24]) + \text{abs}(S0.i[23:16] - S1.i[23:16]) + \text{abs}(S0.i[15:8] - S1.i[15:8]) + \text{abs}(S0.i[7:0] - S1.i[7:0]) + S2.u$. Sum of absolute differences with accumulation, overflow into upper bits is allowed.
- 474 V_SAD_HI_U8: D.u = $(SAD_U8(S0, S1, 0) \ll 16) + S2.u$. Sum of absolute differences with accumulation, overflow is lost.
- 475 V_SAD_UI16: D.u = $\text{abs}(S0.i[31:16] - S1.i[31:16]) + \text{abs}(S0.i[15:0] - S1.i[15:0]) + S2.u$. Word SAD with accumulation.
- 476 V_SAD_U32: D.u = $\text{abs}(S0.i - S1.i) + S2.u$. Dword SAD with accumulation.
- 477 V_CVT_PK_U8_F32: D.u = $\{(\text{flt32_to_uint8}(S0.f) \& 0xff) \ll (8 * S1.u[1:0])\} | (S2.u \& \sim(0xff \ll (8 * S1.u[1:0])))$. Convert floating point value S0 to 8-bit unsigned integer and pack the result into byte S1 of dword S2.
- 478 V_DIV_FIXUP_F32: D.f = Divide fixup and flags -- s0.f = Quotient, s1.f = Denominator, s2.f = Numerator. This opcode generates exceptions resulting from the division operation.
- 479 V_DIV_FIXUP_F64: D.d = Divide fixup and flags -- s0.d = Quotient, s1.d = Denominator, s2.d = Numerator. This opcode generates exceptions resulting from the division operation.
- 480 - 481 reserved.
- 482 V_DIV_FMAS_F32: D.f = Special case divide FMA with scale and flags(s0.f = Quotient, s1.f = Denominator, s2.f = Numerator).
- 483 V_DIV_FMAS_F64: D.d = Special case divide FMA with scale and flags(s0.d = Quotient, s1.d = Denominator, s2.d = Numerator).
- 484 V_MSAD_U8: D.u = Masked Byte SAD with accum_lo(S0.u, S1.u, S2.u).
- 485 V_QSAD_PK_U16_U8: D.u = Quad-Byte SAD with 16-bit packed accum_lo/hi(S0.u[63:0], S1.u[31:0], S2.u[63:0]).
- 486 V_MQSAD_PK_U16_U8: D.u = Masked Quad-Byte SAD with 16-bit packed accum_lo/hi(S0.u[63:0], S1.u[31:0], S2.u[63:0]).
- 487 V_MQSAD_U32_U8: D.u128 = Masked Quad-Byte SAD with 32-bit accum_lo/hi(S0.u[63:0], S1.u[31:0], S2.u[127:0]).
- 488 V_MAD_U64_U32: $\{vcc_out,D.u64\} = S0.u32 * S1.u32 + S2.u64$.
-

Vector Instruction, Three Inputs, One Output (VOP3a)

- 489 V_MAD_I64_I32: {vcc_out,D.i64} = S0.i32 * S1.i32 + S2.i64.
- 490 V_MAD_F16: D.f16 = S0.f16 * S1.f16 + S2.f16 Supports round mode, exception flags, saturation.
- 491 V_MAD_U16: D.u16 = S0.u16 * S1.u16 + S2.u16. Supports saturation (unsigned 16-bit integer domain).
- 492 V_MAD_I16: D.i16 = S0.i16 * S1.i16 + S2.i16. Supports saturation (signed 16-bit integer domain).
- 493 V_PERM_B32: D.u[31:24] = permute({S0.u, S1.u}, S2.u[31:24]);
 D.u[23:16] = permute({S0.u, S1.u}, S2.u[23:16]);
 D.u[15:8] = permute({S0.u, S1.u}, S2.u[15:8]);
 D.u[7:0] = permute({S0.u, S1.u}, S2.u[7:0]);
 byte permute(byte in[8], byte sel) {
 if(sel>=13) then return 0xff;
 elsif(sel==12) then return 0x00;
 elsif(sel==11) then return in[7][7] * 0xff;
 elsif(sel==10) then return in[5][7] * 0xff;
 elsif(sel==9) then return in[3][7] * 0xff;
 elsif(sel==8) then return in[1][7] * 0xff;
 else return in[sel];
 }
 Byte permute.
- 494 V_FMA_F16: D.f16 = S0.f16 * S1.f16 + S2.f16.\n\nFused half precision multiply add.
- 495 V_DIV_FIXUP_F16: sign_out = sign(S1.f16)^sign(S2.f16);
 if (S2.f16 == NAN)
 \tD.f16 = Quiet(S2.f16);
 else if (S1.f16 == NAN)
 \tD.f16 = Quiet(S1.f16);
 else if (S1.f16 == S2.f16 == 0)
 \t# 0/0
 \tD.f16 = pele_nan(0xfe00);
 else if (abs(S1.f16) == abs(S2.f16) == +-INF)
 \t# inf/inf
 \tD.f16 = pele_nan(0xfe00);
 else if (S1.f16 == 0 || abs(S2.f16) == +-INF)
 \t# x/0, or inf/y
 \tD.f16 = sign_out ? -INF : INF;
 else if (abs(S1.f16) == +-INF || S2.f16 == 0)
 \t# x/inf, 0/y
 \tD.f16 = sign_out ? -0 : 0;
 else if ((exp(S2.f16) - exp(S1.f16)) < -150)
 \tD.f16 = sign_out ? -underflow : underflow;
 else if (exp(S1.f16) == 255)
 \tD.f16 = sign_out ? -overflow : overflow;
 else
 \tD.f16 = sign_out ? -abs(S0.f16) : abs(S0.f16).
 Half precision division fixup.
 S0 = Quotient, S1 = Denominator, S3 = Numerator.
 Given a numerator, denominator, and quotient from a divide, this opcode detect sand applies special case numerics, touching up the quotient if necessary. This opcode also generates invalid, denorm, and divide by zero exceptions caused by the division.
-

Vector Instruction, Three Inputs, One Output (VOP3a)

-
- 496 `V_CVT_PKACCUM_U8_F32`: $\text{byte} = S1.u[1:0]$; $\text{bit} = \text{byte} * 8$; $\text{D}.u[\text{bit}+7:\text{bit}] = \text{flt32_to_uint8}(S0.f)$; Pack converted value of $S0.f$ into byte $S1$ of the destination. SQ translates to `V_CVT_PK_U8_F32`. Note: this opcode uses `src_c` to pass destination in as a source.
- 497 – 623 Unused.
- 624 `V_INTERP_P1_F32`: $D = P10 * S + P0$; parameter interpolation.
- 625 `V_INTERP_P2_F32`: $D = P20 * S + D$; parameter interpolation.
- 626 `V_INTERP_MOV_F32`: $D = \{P10, P20, P0\}[S]$; parameter load.
- 627 reserved.
- 628 `V_INTERP_P1LL_F16`: $D.f32 = P10.f16 * S0.f32 + P0.f16$. 'LL' stands for 'two LDS arguments'. `attr_word` selects the high or low half 16 bits of each LDS dword accessed. This opcode is available for 32-bank LDS only. NOTE: In textual representations the I/J VGPR is the first source and the attribute is the second source; however in the VOP3 encoding the attribute is stored in the `src0` field and the VGPR is stored in the `src1` field.
- 629 `V_INTERP_P1LV_F16`: $D.f32 = P10.f16 * S0.f32 + (S2.u32 \gg (\text{attr_word} * 16)).f16$. 'LV' stands for 'One LDS and one VGPR argument'. $S2$ holds two parameters, `attr_word` selects the high or low word of the VGPR for this calculation, as well as the high or low half of the LDS data." Meant for use with 16-bank LDS. NOTE: In textual representations the I/J VGPR is the first source and the attribute is the second source; however in the VOP3 encoding the attribute is stored in the `src0` field and the VGPR is stored in the `src1` field.
- 630 `V_INTERP_P2_F16`: $D.f16 = P20.f16 * S0.f32 + S2.f32$. Final computation. `attr_word` selects LDS high or low 16bits. Used for both 16- and 32-bank LDS. Result is always written to the 16 LSBs of the destination VGPR. NOTE: In textual representations the I/J VGPR is the first source and the attribute is the second source; however in the VOP3 encoding the attribute is stored in the `src0` field and the VGPR is stored in the `src1` field.
- 631 – 639 Unused.
- 640 `V_ADD_F64`: $D.d = S0.d + S1.d$.
- 641 `V_MUL_F64`: $D.d = S0.d * S1.d$.
- 642 `V_MIN_F64`: $D.d = \min(S0.d, S1.d)$.
- 643 `V_MAX_F64`: $D.d = \max(S0.d, S1.d)$.
- 644 `V_LDEXP_F64`: $D.d = \text{pow}(S0.d, S1.i[31:0])$.
- 645 `V_MUL_LO_U32`: $D.u = S0.u * S1.u$.
- 646 `V_MUL_HI_U32`: $D.u = (S0.u * S1.u) \gg 32$.
- 647 `V_MUL_HI_I32`: $D.i = (S0.i * S1.i) \gg 32$.
- 648 `V_LDEXP_F32`: $D.f = \text{pow}(S0.f, S1.i)$.
- 649 `V_READLANE_B32`: Copy one VGPR value to one SGPR. $D = \text{SGPR-dest}$, $S0 = \text{Source Data (VGPR\# or M0(Ids-direct))}$, $S1 = \text{Lane Select (SGPR or M0)}$. Ignores exec mask. Input and output modifiers not supported; this is an untyped operation.
- 650 `V_WRITELANE_B32`: Write value into one VGPR in one lane. $D = \text{VGPR-dest}$, $S0 = \text{Source Data (sgpr, m0, exec or constants)}$, $S1 = \text{Lane Select (SGPR or M0)}$. Ignores exec mask. Input and output modifiers not supported; this is an untyped operation. SQ translates to `V_MOV_B32`.
- 651 `V_BCNT_U32_B32`: $D.u = \text{CountOneBits}(S0.u) + S1.u$. Bit count.
- 652 `V_MBCNT_LO_U32_B32`: $\text{ThreadMask} = (1 \ll \text{ThreadPosition}) - 1$; $\text{D}.u = \text{CountOneBits}(S0.u \& \text{ThreadMask}[31:0]) + S1.u$. Masked bit count, `ThreadPosition` is the position of this thread in the wavefront (in 0..63).
- 653 `V_MBCNT_HI_U32_B32`: $\text{ThreadMask} = (1 \ll \text{ThreadPosition}) - 1$; $\text{D}.u = \text{CountOneBits}(S0.u \& \text{ThreadMask}[63:32]) + S1.u$. Masked bit count, `ThreadPosition` is the position of this thread in the wavefront (in 0..63).
-

Vector Instruction, Three Inputs, One Output (VOP3a)

	654	reserved.
	655	V_LSHLREV_B64: $D.u64 = S1.u64 \ll S0.u[5:0]$.
	656	V_LSHRREV_B64: $D.u64 = S1.u64 \gg S0.u[5:0]$. The vacated bits are set to zero.
	657	V_ASHRREV_I64: $D.u64 = \text{signext}(S1.u64) \gg S0.u[5:0]$. The vacated bits are set to the sign bit of the input value. SQ translates this to an internal SP opcode.
	658	V_TRIG_PREOP_F64: $D.d = \text{Look Up } 2/\text{PI} (S0.d)$ with segment select $S1.u[4:0]$. This operation returns an aligned, double precision segment of $2/\text{PI}$ needed to do range reduction on $S0.d$ (double-precision value). Multiple segments can be specified through $S1.u[4:0]$. Rounding is always round-to-zero. Large inputs ($\text{exp} > 1968$) are scaled to avoid loss of precision through denormalization.
	659	V_BFM_B32: $D.u = ((1 \ll S0.u[4:0]) - 1) \ll S1.u[4:0]$; $S0$ is the bitfield width and $S1$ is the bitfield offset.
	660	V_CVT_PKNORM_I16_F32: $D = \{(\text{snorm})S1.f, (\text{snorm})S0.f\}$.
	661	V_CVT_PKNORM_U16_F32: $D = \{(\text{unorm})S1.f, (\text{unorm})S0.f\}$.
	662	V_CVT_PKRTZ_F16_F32: $D = \{\text{flt32_to_flt16}(S1.f), \text{flt32_to_flt16}(S0.f)\}$, with round-toward-zero regardless of current round mode setting in hardware. This opcode is intended for use with 16-bit compressed exports. See V_CVT_F16_F32 for a version that respects the current rounding mode.
	663	V_CVT_PK_U16_U32: $D = \{\text{uint32_to_uint16}(S1.u), \text{uint32_to_uint16}(S0.u)\}$.
	664	V_CVT_PK_I16_I32: $D = \{\text{int32_to_int16}(S1.i), \text{int32_to_int16}(S0.i)\}$.

ENCODING	[31:26]	enum(6)
	Must be 1 1 0 1 0 0.	

Vector Instruction, Three Inputs, One Output (VOP3a)

SRC0	[40:32]	enum(9)	<p>First operand for instruction.</p> <p>0 – 103 32-bit Scalar General-Purpose Register (SGPR)</p> <p>106 VCC_LO (vcc[31:0]).</p> <p>107 VCC_HI (vcc[63:32]).</p> <p>108 TBA_LO Trap handler base address [31:0]</p> <p>109 TBA_HI Trap handler base address [63:32]).</p> <p>110 TMA_LO Pointer to data in memory used by trap handler.</p> <p>111 TMA_HI Pointer to data in memory used by trap handler.</p> <p>112 TTMP[11:0] Trap handler temporary SGPR [11:0].</p> <p>124 M0. Memory register 0.</p> <p>126 EXEC_LO exec[31:0].</p> <p>127 EXEC_HI exec[63:32].</p> <p>[191 – 128] SRC_[63:0] = 63:0 integer.</p> <p>192 SRC_64_INT = 64 (integer).</p> <p>[208 – 193] SRC_M_[16:1]_INT = [-16:-1] (integer).</p> <p>240 SRC_0_5 = 0.5.</p> <p>241 SRC_M_0_5 = -0.5.</p> <p>242 SRC_1 = 1.0.</p> <p>243 SRC_M_1 = -1.0.</p> <p>244 SRC_2 = 2.0.</p> <p>245 SRC_M_2 = -2.0.</p> <p>246 SRC_4 = 4.0.</p> <p>247 SRC_M_4 = -4.0.</p> <p>251 SRC_VCCZ = vector-condition-code-is-zero.</p> <p>252 SRC_EXECZ = execute-mask-is-zero.</p> <p>253 SRC_SCC = scalar condition code.</p> <p>254 SRC_LDS_DIRECT = use LDS direct to supply 32-bit value (address from M0 register).</p> <p>256 - 511 VGPR 0 to 255.</p> <p>All other values are reserved.</p>
SRC1	[49:41]	enum(9)	<p>Second operand for instruction. Same format as SRC0.</p>
SRC2	[58:50]	enum(9)	<p>Third operand for instruction. Same format as SRC0.</p>
OMOD	[60:59]	enum(2)	<p>Output modifier for instruction. Applied before clamping.</p> <p>0 : No modification.</p> <p>1 : Multiply output by 2.0.</p> <p>2 : Multiply output by 4.0.</p> <p>3 : Divide output by 2.0.</p>
NEG	[63:61]	enum(3)	<p>If NEG[N] is set, take the floating-point negation of the N'th input operand. This is applied after absolute value.</p>

Vector Instruction, Three Inputs, One Vector Output and One Scalar Output (VOP3b)

<i>Format</i>	VOP3b		
<i>Description</i>	Vector instruction taking three inputs and producing two outputs.		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	VDST	[7:0]	enum(8) Destination for instruction in the Vector General-Purpose Registers (VGPR [255:0]).
	SDST	[14:8]	enum(7) 0 – 103 32-bit Scalar General-Purpose Register (SGPR) 106 VCC_LO (vcc[31:0]). 107 VCC_HI (vcc[63:32]). 108 TBA_LO Trap handler base address [31:0] 109 TBA_HI Trap handler base address [63:32]). 110 TMA_LO Pointer to data in memory used by trap handler. 111 TMA_HI Pointer to data in memory used by trap handler. 113 – 112 TTMP[11:0] Trap handler temporary SGPR [11:0]. All other values are reserved.
	CLAMP	15	enum(1) If set, clamp output to [0.0, 1.0]. Applied after output modifier.
	OP	[25:16]	enum(10) Instructions that use this format: 281 V_ADD_U32: D.u = S0.u + S1.u; VCC=carry-out (VOP3:sgpr=carry-out). 282 V_SUB_U32: D.u = S0.u - S1.u; VCC=carry-out (VOP3:sgpr=carry-out). 283 V_SUBREV_U32: D.u = S1.u - S0.u; VCC=carry-out (VOP3:sgpr=carry-out). 284 V_ADDC_U32: D.u = S0.u + S1.u + VCC; VCC=carry-out (VOP3:sgpr=carry-out, S2.u=carry-in). 285 V_SUBB_U32: D.u = S0.u - S1.u - VCC; VCC=carry-out (VOP3:sgpr=carry-out, S2.u=carry-in). 286 V_SUBBREV_U32: D.u = S1.u - S0.u - VCC; VCC=carry-out (VOP3:sgpr=carry-out, S2.u=carry-in). 480 V_DIV_SCALE_F32: {vcc,D.f} = Divide preop and flags -- s0.f = Quotient, s1.f = Denominator, s2.f = Numerator -- s0 must equal s1 or s2. Given a numerator and denominator, this opcode appropriately scales inputs for division to avoid subnormal terms during Newton-Raphson correction algorithm. This opcode produces a VCC flag for post-scale of quotient 481 V_DIV_SCALE_F64: {vcc,D.d} = Divide preop and flags -- s0.d = Quotient, s1.d = Denominator, s2.d = Numerator -- s0 must equal s1 or s2. Given a numerator and denominator, this opcode appropriately scales inputs for division to avoid subnormal terms during Newton-Raphson correction algorithm. This opcode produces a VCC flag for post-scale of quotient.

Vector Instruction, Three Inputs, One Vector Output and One Scalar Output (VOP3b)

Group 1 and Group 2, below are the VOPC opcodes when VOP3 encoding is required.

Sixteen Compare Operations (OP16)

Compare Opcode

<u>Operation</u>	<u>Offset</u>	<u>Description</u>
F	0	D.u = 0
LT	1	D.u = (S0 < S1)
EQ	2	D.u = (S0 == S1)
LE	3	D.u = (S0 <= S1)
GT	4	D.u = (S0 > S1)
LG	5	D.u = (S0 <> S1)
GE	6	D.u = (S0 >= S1)
O	7	D.u = (!isNaN(S0) && !isNaN(S1))
U	8	D.u = (!isNaN(S0) !isNaN(S1))
NGE	9	D.u = !(S0 >= S1)
NLG	10	D.u = !(S0 <> S1)
NGT	11	D.u = !(S0 > S1)
NLE	12	D.u = !(S0 <= S1)
NEQ	13	D.u = !(S0 == S1)
NLT	14	D.u = !(S0 < S1)
TRU	15	D.u = 1

Eight Compare Operations (OP8)

Compare Opcode

<u>Operation</u>	<u>Offset</u>	<u>Description</u>
F	0	D.u = 0
LT	1	D.u = (S0 < S1)
EQ	2	D.u = (S0 == S1)
LE	3	D.u = (S0 <= S1)
GT	4	D.u = (S0 > S1)
LG	5	D.u = (S0 <> S1)
GE	6	D.u = (S0 >= S1)
TRU	7	D.u = 1

Vector Instruction, Three Inputs, One Vector Output and One Scalar Output (VOP3b)

Single Vector Compare Operations

<u>Opcode Family</u>	<u>Opcode</u>	
	<u>Base</u>	<u>Description</u>
Reserved	0x00	
CMP_CLASS_F32	0x10	none
CMPX_CLASS_F32	0x11	none
CMP_CLASS_F64	0x12	none
CMPX_CLASS_F64	0x13	none
CMP_CLASS_F16	0x14	none
CMPX_CLASS_F16	0x15	none
Reserved	0x16-0x1F	
CMP_F16	0x20	OP16
CMPX_F16	0x30	OP16
CMP_F32	0x40	OP16
CMPX_F32	0x50	OP16
CMP_F64	0x60	OP16
CMPX_F64	0x70	OP16
Reserved	0x80-0x98	
CMP_I16	0xA0	OP8
CMP_U16	0xA8	OP8
CMPX_I16	0xB0	OP8
CMPX_U16	0xB8	OP8
CMP_I32	0xC0	OP8
CMP_U32	0xC8	OP8
CMPX_I32	0xD0	OP8
CMPX_U32	0xD8	OP8
CMP_I64	0xE0	OP8
CMP_U64	0xE8	OP8
CMPX_I64	0xF0	OP8
CMPX_U64	0xF8	OP8
ENCODING	[31:26]	enum(6)
	Must be 1 1 0 1 0 0.	

Vector Instruction, Three Inputs, One Vector Output and One Scalar Output (VOP3b)

SRC0	[40:32]	enum(9)
	First operand for instruction. 0 – 103 32-bit Scalar General-Purpose Register (SGPR) 106 VCC_LO (vcc[31:0]). 107 VCC_HI (vcc[63:32]). 108 TBA_LO Trap handler base address [31:0] 109 TBA_HI Trap handler base address [63:32]). 110 TMA_LO Pointer to data in memory used by trap handler. 111 TMA_HI Pointer to data in memory used by trap handler. 113 – 112 TTMP[11:0] Trap handler temporary SGPR [11:0]. 124 M0. Memory register 0. 126 EXEC_LO exec[31:0]. 127 EXEC_HI exec[63:32]. [191 – 128] SRC_[63:0] = 63:0 integer. 192 SRC_64_INT = 64 (integer). [208 – 193] SRC_M_[16:1]_INT = [-16:-1] (integer). 240 SRC_0_5 = 0.5. 241 SRC_M_0_5 = -0.5. 242 SRC_1 = 1.0. 243 SRC_M_1 = -1.0. 244 SRC_2 = 2.0. 245 SRC_M_2 = -2.0. 246 SRC_4 = 4.0. 247 SRC_M_4 = -4.0. 251 SRC_VCCZ = vector-condition-code-is-zero. 252 SRC_EXECZ = execute-mask-is-zero. 253 SRC_SCC = scalar condition code. 254 SRC_LDS_DIRECT = use LDS direct to supply 32-bit value (address from M0 register). 256 – 511 VGPR 0 to 255. All other values are reserved.	
SRC1	[49:41]	enum(9)
	Second operand for instruction.	
SRC2	[58:50]	enum(9)
	Third operand for instruction.	
OMOD	[60:59]	enum(2)
	Output modifier for instruction. Applied before clamping.	
NEG	[63:61]	enum(3)
	If NEG[N] is set, take the floating-point negation of the N'th input operand. This is applied after absolute value.	

Vector Instruction with Sub-Dword Addressing

Format	VOP_SDWA		
Description	Second Dword for VOP1/VOP2/VOPC instructions for specifying sub-dword addressing flags.		
Opcod	Field Name	Bits	Format
	SRCO	[7:0]	enum(8) Vector General-Purpose Registers (VGPR) address for src0.
	DST_SEL	[10:8]	enum(3) Destination data select. 0 SDWA_BYTE_0: Select data[7:0]. S DWA_BYTE_1: Select data[15:8]. 2 SDWA_BYTE_2: Select data[23:16]. 3 SDWA_BYTE_3: Select data[31:24]. 4 SDWA_WORD_0: Select data[15:0]. 5 SDWA_WORD_1: Select data[31:16]. 6 SDWA_DWORD: Select data[31:0]. 7 reserved,
	DST_UNUSED	[12:11]	enum(2) Format for unused destination bits. 0 SDWA_UNUSED_PAD: Pad all unused bits with 0. 1 SDWA_UNUSED_SEXT: Sign-extend upper bits; pad lower bits with 0. 2 SDWA_UNUSED_PRESERVE : Select data[31:0]. 3 reserved.
	CLAMP	13	enum(1) If true: clamp output.
	SRCO_SEL	[18:16]	enum(3) Source data select for src0. 0 SDWA_BYTE_0: Select data[7:0]. S DWA_BYTE_1: Select data[15:8]. 2 SDWA_BYTE_2: Select data[23:16]. 3 SDWA_BYTE_3: Select data[31:24]. 4 SDWA_WORD_0: Select data[15:0]. 5 SDWA_WORD_1: Select data[31:16]. 6 SDWA_DWORD: Select data[31:0]. 7 reserved.
	SRCO_SEXT	19	enum(1) If true, sign-extend data for src0. If false, zero-extend.
	SRCO_NEG	20	enum(1) If true, take the floating-point negation of src0.
	SRCO_ABS	21	enum(1) If true, take the floating-point absolute value of src0.
	Reserved	[23:22]	Reserved.

Vector Instruction with Sub-Dword Addressing

SRC1_SEL	[26:24]	enum(3)
	Source data select for src1.	
	0	SDWA_BYTE_0: Select data[7:0].
	S	DWA_BYTE_1: Select data[15:8].
	2	SDWA_BYTE_2: Select data[23:16].
	3	SDWA_BYTE_3: Select data[31:24].
	4	SDWA_WORD_0: Select data[15:0].
	5	SDWA_WORD_1: Select data[31:16].
	6	SDWA_DWORD: Select data[31:0].
	7	reserved.

SRC1_SEXT	27	enum(1)
	If true, sign-extend data for src1. If false, zero-extend.	

SRC1_NEG	28	enum(1)
	If true, take the floating-point negation of src1.	

SRC1_ABS	29	enum(1)
	If true, take the floating-point absolute value of src1.	

reserved	[31:30]	enum(2)
	Reserved.	

Vector Instruction with Data Parallel Primitives

Format **VOP_DPP**

Description Second Dword for VOP1/VOP2/VOPC instructions for specifying data-parallel primitives.

Opcode	Field Name	Bits	Format
	SRCO	[7:0]	enum(8) Vector General-Purpose Registers (VGPR) address for src0.
	DPP_CTRL	[16:8]	enum(9) Data-parallel primitive control.

dpp_cntl Enumeration	Hex Value	Function	Description
DPP_QUAD_PERM{00:FF}* FF}	000-0FF	$\text{pix}[n].\text{srca} = \text{pix}[(n \& 0x3c) + \text{dpp_cntl}[n\%4*2+1 : n\%4*2]].\text{srca}$	Full permute of four threads.
DPP_UNUSED	100	Undefined	Reserved.
DPP_ROW_SL{1:15}* 15}	101-10F	if $((n \& 0xf) < (16 - \text{cntl}[3:0]))$ $\text{pix}[n].\text{srca} = \text{pix}[n + \text{cntl}[3:0]].\text{srca}$ else use bound_cntl	Row shift right by 1-15 threads.
DPP_ROW_SR{1:15}* 15}	111-11F	if $((n \& 0xf) \geq \text{cntl}[3:0])$ $\text{pix}[n].\text{srca} = \text{pix}[n - \text{cntl}[3:0]].\text{srca}$ else use bound_cntl	Row shift right by 1-15 threads.
DPP_ROW_RR{1:15}* 15}	121-12F	if $((n \& 0xf) \geq \text{cntl}[3:0])$ $\text{pix}[n].\text{srca} = \text{pix}[n - \text{cntl}[3:0]].\text{srca}$ else $\text{pix}[n].\text{srca} = \text{pix}[n + 16 - \text{cntl}[3:0]].\text{srca}$	Row rotate right by 1-15 threads.
DPP_WF_SL1* 1	130	if $(n < 63)$ $\text{pix}[n].\text{srca} = \text{pix}[n+1].\text{srca}$ else use bound_cntl	Wavefront left shift by 1 thread.
DPP_WF_RL1* 1	134	if $(n < 63)$ $\text{pix}[n].\text{srca} = \text{pix}[n+1].\text{srca}$ else $\text{pix}[n].\text{srca} = \text{pix}[0].\text{srca}$	Wavefront left rotate by 1 thread.
DPP_WF_SR1* 1	138	if $(n > 0)$ $\text{pix}[n].\text{srca} = \text{pix}[n-1].\text{srca}$ else use bound_cntl	Wavefront right shift by 1 thread.
DPP_WF_RR1* 1	13C	if $(n > 0)$ $\text{pix}[n].\text{srca} = \text{pix}[n-1].\text{srca}$ else $\text{pix}[n].\text{srca} = \text{pix}[63].\text{srca}$	Wavefront right rotate by 1 thread.
DPP_ROW_MIRROR* 1	140	$\text{pix}[n].\text{srca} = \text{pix}[15-(n \& f)].\text{srca}$	Mirror threads within row..
DPP_ROW_HALF_MIRROR* 1	141	$\text{pix}[n].\text{srca} = \text{pix}[7-(n \& 7)].\text{srca}$	Mirror threads within 1/2 row (8 threads).
DPP_ROW_BCAST15* 15	142	if $(n > 15)$ $\text{pix}[n].\text{srca} = \text{pix}[n \& 0x30 - 1].\text{srca}$	Broadcast 15 th thread of each row to next row.
DPP_ROW_BCAST31* 31	143	if $(n > 31)$ $\text{pix}[n].\text{srca} = \text{pix}[n \& 0x20 - 1].\text{srca}$	Broadcast thread 31 to rows 2 and 3.

reserved [18:17]

Reserved.

Vector Instruction with Data Parallel Primitives

BOUND_CTRL	19	enum(1)	Specifies behavior when shared data is invalid. 0 Set write-enable to zero. 1 Set source data to zero.
SRC0_NEG	20	enum(1)	If true, take the floating-point negation of src0.
SRC0_ABS	21	enum(1)	If true, take the floating-point absolute value of src0.
SRC1_NEG	22	enum(1)	If true, take the floating-point negation of src1.
SRC1_ABS	23	enum(1)	If true, take the floating-point absolute value of src1.
BANK_MASK	[27:24]	enum(4)	Bank enable mask. If bit N is set, then lanes[4*N:4*N+3, 4*N+16:4*N+19, 4*N+32:4*N+35, 4*N+48:4*N+51] are enabled.
ROW_MASK	[31:28]	enum(4)	Row enable mask. If bit N is set, then lanes[16*N:16*N+15] are enabled.

13.4 Vector Parameter Interpolation Instruction

Interpolation Instruction

<i>Format</i>	VINTRP		
<i>Description</i>	Interpolate data for the pixel shader.		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	VSRC	[7:0]	enum(8) Vector general-purpose registers (VGPR) containing the i/j coordinate by which to multiply one of the parameter components.
	ATTRCHAN	[9:8]	enum(2) Attribute component to interpolate. See Section 10.1 on page 10-1.
	ATTR	[15:10]	int(6) Attribute to interpolate.
	OP	[17:16]	enum(2) 0 V_INTERP_P1_F32: $D = P10 * S + P0$; parameter interpolation. 1 V_INTERP_P2_F32: $D = P20 * S + D$; parameter interpolation. 2 V_INTERP_MOV_F32: $D = \{P10, P20, P0\}[S]$; parameter load. 3 Reserved.
	VDST	[25:18]	enum(8) Vector general-purpose registers VGPR [255:0] to which results are written, and, optionally, from which they are read when accumulating results.
	ENCODING	[31:26]	enum(6) Must be 1 1 0 0 1 0.

13.5 LDS/GDS Instruction

Data Share Instruction

Format	DS		
Description	.Local and global data share instructions.		
Opcode	Field Name	Bits	Format
	OFFSET0	[7:0]	int(8) Unsigned byte offset added to the address supplied by the ADDR VGPR.
	OFFSET1	[15:8]	int(8) Unsigned byte offset added to the address supplied by the ADDR VGPR.
	GDS	16	enum(1) 0 = LDS; 1 = GDS.
	OP	[24:17]	enum(8) uint = unsigned integer; int = signed integer. 00 DS_ADD_U32: DS[A] = DS[A] + D0; uint add. 01 DS_SUB_U32: DS[A] = DS[A] - D0; uint subtract. 02 DS_RSUB_U32: DS[A] = D0 - DS[A]; uint reverse subtract. 03 DS_INC_U32: DS[A] = (DS[A] >= D0 ? 0 : DS[A] + 1); uint increment. 04 DS_DEC_U32: DS[A] = (DS[A] == 0 DS[A] > D0 ? D0 : DS[A] - 1); uint decrement. 05 DS_MIN_I32: DS[A] = min(DS[A], D0); int min. 06 DS_MAX_I32: DS[A] = max(DS[A], D0); int max. 07 DS_MIN_U32: DS[A] = min(DS[A], D0); uint min. 08 DS_MAX_U32: DS[A] = max(DS[A], D0); uint max. 09 DS_AND_B32: DS[A] = DS[A] & D0; Dword AND. 10 DS_OR_B32: DS[A] = DS[A] D0; Dword OR. 11 DS_XOR_B32: DS[A] = DS[A] ^ D0; Dword XOR. 12 DS_MSKOR_B32: DS[A] = (DS[A] & ~D0) D1; masked Dword OR. 13 DS_WRITE_B32: DS[A] = D0; write a Dword. 14 DS_WRITE2_B32: DS[ADDR+offset0*4] = D0; DS[ADDR+offset1*4] = D1; write 2 Dwords. 15 DS_WRITE2ST64_B32: DS[ADDR+offset0*4*64] = D0; DS[ADDR+offset1*4*64] = D1; write 2 Dwords. 16 DS_CMPST_B32: DS[A] = (DS[A] == D0 ? D1 : DS[A]); compare store. 17 DS_CMPST_F32: DS[A] = (DS[A] == D0 ? D1 : DS[A]); compare store with float rules. 18 DS_MIN_F32: DS[A] = (DS[A] < D1) ? D0 : DS[A]; float compare swap (handles NaN/INF/denorm). 19 DS_MAX_F32: DS[A] = (D0 > DS[A]) ? D0 : DS[A]; float, handles NaN/INF/denorm. 20 DS_NOP Do nothing. 21 DS_ADD_F32: DS[A] = DS[A] + D0; float add. 30 DS_WRITE_B8: DS[A] = D0[7:0]; byte write. 31 DS_WRITE_B16: DS[A] = D0[15:0]; short write. 32 DS_ADD_RTN_U32: Uint add. 33 DS_SUB_RTN_U32: Uint subtract. 34 DS_RSUB_RTN_U32: Uint reverse subtract. 35 DS_INC_RTN_U32: Uint increment.

Data Share Instruction

36	DS_DEC_RTN_U32: Uint decrement.
37	DS_MIN_RTN_I32: Int min.
38	DS_MAX_RTN_I32: Int max.
39	DS_MIN_RTN_U32: Uint min.
40	DS_MAX_RTN_U32: Uint max.
41	DS_AND_RTN_B32: Dword AND.
42	DS_OR_RTN_B32: Dword OR.
43	DS_XOR_RTN_B32: Dword XOR.
44	DS_MSKOR_RTN_B32: Masked Dword OR.
45	DS_WRXCHG_RTN_B32: Write exchange. Offset = {offset1,offset0}. A = ADDR+offset. D=DS[Addr]. DS[Addr]=D0.
46	DS_WRXCHG2_RTN_B32: Write exchange 2 separate Dwords.
47	DS_WRXCHG2ST64_RTN_B32: Write exchange 2 Dwords, stride 64.
48	DS_CMPST_RTN_B32: Compare store.
49	DS_CMPST_RTN_F32: Compare store with float rules.
50	DS_MIN_RTN_F32: DS[A] = (D0 < DS[A]) ? D0 : DS[A]; float compare swap (handles NaN/INF/denorm).
51	DS_MAX_RTN_F32: DS[A] = (D0 > DS[A]) ? D0 : DS[A]; float, handles NaN/INF/denorm .
52	DS_WRAP_RTN_B32: DS[A] = (DS[A] >= D0) ? DS[A] - D0 : DS[A] + D1.
53	reserved.
61	DS_SWIZZLE_B32: R = swizzle(Data(VGPR), offset1:offset0). Dword swizzle. no data is written to LDS.
54	DS_READ_B32: R = DS[A]; Dword read.
55	DS_READ2_B32: R = DS[ADDR+offset0*4], R+1 = DS[ADDR+offset1*4]. Read two Dwords.
56	DS_READ2ST64_B32: R = DS[ADDR+offset0*4*64], R+1 = DS[ADDR+offset1*4*64]. Read two Dwords.
57	DS_READ_I8: R = signext(DS[A][7:0]); signed byte read.
58	DS_READ_U8: R = {24'h0,DS[A][7:0]}; unsigned byte read.
59	DS_READ_I16: R = signext(DS[A][15:0]); signed short read.
60	DS_READ_U16: R = {16'h0,DS[A][15:0]}; unsigned short read.
62	DS_PERMUTE_B32: Forward permute. Does not write any LDS memory. LDS[dst] = src0 returnVal = LDS[thread_id] Where "thread_id" is 0..63.
63	DS_BPERMUTE_B32: Backward permute. Does not actually write any LDS mem- ory. LDS[thread_id] = src0 Where "thread_id" is 0..63. returnVal = LDS[dst]
64	DS_ADD_U64: Uint add.
65	DS_SUB_U64: Uint subtract.
66	DS_RSUB_U64: Uint reverse subtract.
67	DS_INC_U64: Uint increment.
68	DS_DEC_U64: Uint decrement.
69	DS_MIN_I64: Int min.
70	DS_MAX_I64: Int max.
71	DS_MIN_U64: Uint min.
72	DS_MAX_U64: Uint max.
73	DS_AND_B64: Dword AND.
74	DS_OR_B64: Dword OR.
75	DS_XOR_B64: Dword XOR.

Data Share Instruction

-
- 76 DS_MSKOR_B64: Masked Dword XOR.
 - 77 DS_WRITE_B64: Write.
 - 78 DS_WRITE2_B64: DS[ADDR+offset0*8] = D0; DS[ADDR+offset1*8] = D1; write 2 Dwords.
 - 79 DS_WRITE2ST64_B64: DS[ADDR+offset0*8*64] = D0; DS[ADDR+offset1*8*64] = D1; write 2 Dwords.
 - 80 DS_CMPST_B64: Compare store.
 - 81 DS_CMPST_F64: Compare store with float rules.
 - 82 DS_MIN_F64: DS[A] = (D0 < DS[A]) ? D0 : DS[A]; float, handles NaN/INF/denorm.
 - 83 DS_MAX_F64: DS[A] = (D0 > DS[A]) ? D0 : DS[A]; float, handles NaN/INF/denorm.
 - 96 DS_ADD_RTN_U64: Uint add.
 - 97 DS_SUB_RTN_U64: Uint subtract.
 - 98 DS_RSUB_RTN_U64: Uint reverse subtract.
 - 99 DS_INC_RTN_U64: Uint increment.
 - 100 DS_DEC_RTN_U64: Uint decrement.
 - 101 DS_MIN_RTN_I64: Int min.
 - 102 DS_MAX_RTN_I64: Int max.
 - 103 DS_MIN_RTN_U64: Uint min.
 - 104 DS_MAX_RTN_U64: Uint max.
 - 105 DS_AND_RTN_B64: Dword AND.
 - 106 DS_OR_RTN_B64: Dword OR.
 - 107 DS_XOR_RTN_B64: Dword XOR.
 - 108 DS_MSKOR_RTN_B64: Masked Dword XOR.
 - 109 DS_WRXCHG_RTN_B64: Write exchange.
 - 110 DS_WRXCHG2_RTN_B64: Write exchange relative.
 - 111 DS_WRXCHG2ST64_RTN_B64: Write exchange 2 Dwords.
 - 112 DS_CMPST_RTN_B64: Compare store.
 - 113 DS_CMPST_RTN_F64: Compare store with float rules.
 - 114 DS_MIN_RTN_F64: DS[A] = (D0 < DS[A]) ? D0 : DS[A]; float, handles NaN/INF/denorm.
 - 115 DS_MAX_RTN_F64: DS[A] = (D0 > DS[A]) ? D0 : DS[A]; float, handles NaN/INF/denorm.
 - 118 DS_READ_B64: Dword read.
 - 119 DS_READ2_B64: R = DS[ADDR+offset0*8], R+1 = DS[ADDR+offset1*8]. Read 2 Dwords
 - 120 DS_READ2ST64_B64: R = DS[ADDR+offset0*8*64], R+1 = DS[ADDR+offset1*8*64]. Read 2 Dwords.
 - 126 DS_CONDXCHG32_RTN_B64: Conditional write exchange.
 - 128 DS_ADD_SRC2_U32: B = A + 4*(offset1[7] ? {A[31],A[31:17]} : {offset1[6],offset1[6:0],offset0}). DS[A] = DS[A] + DS[B]; uint add.
 - 129 DS_SUB_SRC2_U32: B = A + 4*(offset1[7] ? {A[31],A[31:17]} : {offset1[6],offset1[6:0],offset0}). DS[A] = DS[A] - DS[B]; uint subtract.
 - 130 DS_RSUB_SRC2_U32: B = A + 4*(offset1[7] ? {A[31],A[31:17]} : {offset1[6],offset1[6:0],offset0}). DS[A] = DS[B] - DS[A]; uint reverse subtract.
 - 131 DS_INC_SRC2_U32: B = A + 4*(offset1[7] ? {A[31],A[31:17]} : {offset1[6],offset1[6:0],offset0}). DS[A] = (DS[A] >= DS[B] ? 0 : DS[A] + 1); uint increment.
-

Data Share Instruction

-
- 132 DS_DEC_SRC2_U32: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = (DS[A] == 0 || DS[A] > DS[B] ? DS[B] : DS[A] - 1); uint decrement.
- 133 DS_MIN_SRC2_I32: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = min(DS[A], DS[B]); int min.
- 134 DS_MAX_SRC2_I32: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = max(DS[A], DS[B]); int max.
- 135 DS_MIN_SRC2_U32: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = min(DS[A], DS[B]); uint min.
- 136 DS_MAX_SRC2_U32: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = max(DS[A], DS[B]); uint maxw.
- 137 DS_AND_SRC2_B32: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = DS[A] & DS[B]; Dword AND.
- 138 DS_OR_SRC2_B32: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = DS[A] | DS[B]; Dword OR.
- 139 DS_XOR_SRC2_B32: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = DS[A] ^ DS[B]; Dword XOR.
- 141 DS_WRITE_SRC2_B32: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = DS[B]; write Dword.
- 146 DS_MIN_SRC2_F32: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = (DS[B] < DS[A] ? DS[B] : DS[A]; float, handles NaN/INF/denorm.
- 147 DS_MAX_SRC2_F32: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$. DS[A] = (DS[B] > DS[A] ? DS[B] : DS[A]; float, handles NaN/INF/denorm.
- 152 DS_GWS_SEMA_RELEASE_ALL: GDS Only. Release all wavefronts waiting on this semaphore. ResourceID is in offset[4:0].
- 153 DS_GWS_INIT: GDS only.
- 154 DS_GWS_SEMA_V: GDS only.
- 155 DS_GWS_SEMA_BR: GDS only.
- 156 DS_GWS_SEMA_P: GDS only.
- 157 DS_GWS_BARRIER: GDS only.
- 189 DS_CONSUME: Consume entries from a buffer.
- 190 DS_APPEND: Append one or more entries to a buffer.
- 191 DS_ORDERED_COUNT: Increment an append counter. Operation is done in order of wavefront creation.
- 192 DS_ADD_SRC2_U64: Uint add.
- 193 DS_SUB_SRC2_U64: Uint subtract.
- 194 DS_RSUB_SRC2_U64: Uint reverse subtract.
- 195 DS_INC_SRC2_U64: Uint increment.
- 196 DS_DEC_SRC2_U64: Uint decrement.
- 197 DS_MIN_SRC2_I64: Int min.
- 198 DS_MAX_SRC2_I64: Int max.
-

Data Share Instruction

- 199 DS_MIN_SRC2_U64: Uint min.
 - 200 DS_MAX_SRC2_U64: Uint max.
 - 201 DS_AND_SRC2_B64: Dword AND.
 - 202 DS_OR_SRC2_B64: Dword OR.
 - 203 DS_XOR_SRC2_B64: Dword XOR.
 - 205 DS_WRITE_SRC2_B64: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$; DS[A] = DS[B]; write Qword.
 - 210 DS_MIN_SRC2_F64: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$; [A] = (D0 < DS[A]) ? D0 : DS[A]; float, handles NaN/INF/denorm.
 - 211 DS_MAX_SRC2_F64: $B = A + 4 * (\text{offset1}[7] ? \{A[31], A[31:17]\} : \{\text{offset1}[6], \text{offset1}[6:0], \text{offset0}\})$; [A] = (D0 > DS[A]) ? D0 : DS[A]; float, handles NaN/INF/denorm.
 - 222 DS_WRITE_B96: {DS[A+2], DS[A+1], DS[A]} = D0[95:0]; tri-Dword write.
 - 223 DS_WRITE_B128: {DS[A+3], DS[A+2], DS[A+1], DS[A]} = D0[127:0]; qword write.
 - 253 DS_CONDXCHG32_RTN_B128: Conditional write exchange.
 - 254 DS_READ_B96: Tri-Dword read.
 - 255 DS_READ_B128: Qword read.
- All other values are reserved.

reserved	25	Reserved.
<hr/>		
ENCODING	[31:26]	enum(6) Must be 1 1 0 1 1 0.
<hr/>		
ADDR	[39:32]	enum(8) Source LDS address VGPR 0 - 255.
<hr/>		
DATA0	[47:40]	enum(8) Source data0 VGPR 0 - 255.
<hr/>		
DATA1	[55:48]	enum(8) Source data1 VGPR 0 - 255.
<hr/>		
VDST	[63:56]	enum(8) Destination VGPR 0 - 255.

13.6 Vector Memory Buffer Instructions

Untyped Vector Memory Buffer Operation

<i>Format</i>	MUBUF		
<i>Description</i>	Untyped memory buffer operation. First word with LDS, second word non-LDS.		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	OFFSET	[11:0]	int(12)
			Unsigned byte offset.
	OFFEN	12	enum(1)
			If set, send VADDR as an offset. If clear, use zero instead of an offset from a VGPR.
	IDXEN	13	enum(1)
			If set, send VADDR as an index. If clear, treat the index as zero.
	GLC	14	enum(1)
			If set, operation is globally coherent.
reserved	15		Reserved.
LDS	16	enum(1)	
			If set, data is read from/written to LDS memory. If unset, data is read from/written to a VGPR.
SLC	17	enum(1)	
			System Level Coherent.

Untyped Vector Memory Buffer Operation

OP	[24:18]	enum(8)
0 – 7 Types buffer loads/stores.		
0	BUFFER_LOAD_FORMAT_X:	Untyped buffer load one Dword with format conversion.
1	BUFFER_LOAD_FORMAT_XY:	Untyped buffer load two Dwords with format conversion.
2	BUFFER_LOAD_FORMAT_XYZ:	Untyped buffer load three Dwords with format conversion.
3	BUFFER_LOAD_FORMAT_XYZW:	Untyped buffer load four Dwords with format conversion.
4	BUFFER_STORE_FORMAT_X:	Untyped buffer store one Dword with format conversion.
5	BUFFER_STORE_FORMAT_XY:	Untyped buffer store two Dwords with format conversion.
6	BUFFER_STORE_FORMAT_XYZ:	Untyped buffer store three Dwords with format conversion.
7	BUFFER_STORE_FORMAT_XYZW:	Untyped buffer store four Dwords with format conversion.
8	BUFFER_LOAD_FORMAT_D16_X:	Untyped buffer load 1 dword with format conversion.
9	BUFFER_LOAD_FORMAT_D16_XY:	Untyped buffer load 2 dwords with format conversion.
10	BUFFER_LOAD_FORMAT_D16_XYZ:	Untyped buffer load 3 dwords with format conversion.
11	BUFFER_LOAD_FORMAT_D16_XYZW:	Untyped buffer load 4 dwords with format conversion.
12	BUFFER_STORE_FORMAT_D16_X:	Untyped buffer store 1 dword with format conversion.
13	BUFFER_STORE_FORMAT_D16_XY:	Untyped buffer store 2 dwords with format conversion.
14	BUFFER_STORE_FORMAT_D16_XYZ:	Untyped buffer store 3 dwords with format conversion.
15	BUFFER_STORE_FORMAT_D16_XYZW:	Untyped buffer store 4 dwords with format conversion.
16	BUFFER_LOAD_UBYTE:	Untyped buffer load unsigned byte (zero extend to VGPR destination).
17	BUFFER_LOAD_SBYTE:	Untyped buffer load signed byte (sign extend to VGPR destination).
18	BUFFER_LOAD_USHORT:	Untyped buffer load unsigned short (zero extend to VGPR destination).
19	BUFFER_LOAD_SSHORT:	Untyped buffer load signed short (sign extend to VGPR destination).
20	BUFFER_LOAD_DWORD:	Untyped buffer load Dword.
21	BUFFER_LOAD_DWORDX2:	Untyped buffer load 2 Dwords.
22	BUFFER_LOAD_DWORDX3:	Untyped buffer load 3 Dwords.
23	BUFFER_LOAD_DWORDX4:	Untyped buffer load 4 Dwords.

Untyped Vector Memory Buffer Operation

24 – 31 Untyped buffer stores.

- 24 BUFFER_STORE_BYTE: Untyped buffer store byte.
 25 reserved.
 26 BUFFER_STORE_SHORT: Untyped buffer store short.
 27 reserved.
 28 BUFFER_STORE_DWORD: Untyped buffer store Dword.
 29 BUFFER_STORE_DWORDX2: Untyped buffer store 2 Dwords.
 30 BUFFER_STORE_DWORDX3: Untyped buffer store 3 Dwords.
 31 BUFFER_STORE_DWORDX4: Untyped buffer store 4 Dwords.
 32 – 60 reserved.
 61 BUFFER_STORE_LDS_DWORD: Store one Dword from LDS memory to system memory without using VGPRs.
 62 BUFFER_WBINVL1: Write back and invalidate the shader L1. Always returns ACK to shader.
 63 BUFFER_WBINVL1_VOL: Write back and invalidate the shader L1 only for lines that are marked volatile. Always returns ACK to shader.
-

64 – 79 Atomic (single).

- 64 BUFFER_ATOMIC_SWAP: 32b:tmp = MEM[ADDR];\nMEM[ADDR] = DATA; RETURN_DATA = tmp.
 65 BUFFER_ATOMIC_CMPSWAP: 32b:tmp = MEM[ADDR];src = DATA[0];cmp = DATA[1]; MEM[ADDR] = (tmp == cmp) ? src : tmp; RETURN_DATA[0] = tmp.
 66 BUFFER_ATOMIC_ADD: 32b:tmp = MEM[ADDR];MEM[ADDR] += DATA; RETURN_DATA = tmp.
 67 BUFFER_ATOMIC_SUB: 32b:tmp = MEM[ADDR];MEM[ADDR] -= DATA; RETURN_DATA = tmp.
 68 BUFFER_ATOMIC_SMIN: 32b:tmp = MEM[ADDR];MEM[ADDR] = (DATA < tmp) ? DATA : tmp (signed compare); RETURN_DATA = tmp.
 69 BUFFER_ATOMIC_UMIN: 32b:tmp = MEM[ADDR];MEM[ADDR] = (DATA < tmp) ? DATA : tmp (unsigned compare); RETURN_DATA = tmp.
 70 BUFFER_ATOMIC_SMAX: 32b:tmp = MEM[ADDR];MEM[ADDR] = (DATA > tmp) ? DATA : tmp (signed compare); RETURN_DATA = tmp.
 71 BUFFER_ATOMIC_UMAX: 32b:tmp = MEM[ADDR];MEM[ADDR] = (DATA > tmp) ? DATA : tmp (unsigned compare); RETURN_DATA = tmp.
 72 BUFFER_ATOMIC_AND: 32b:tmp = MEM[ADDR];MEM[ADDR] &= DATA; RETURN_DATA = tmp.
 73 BUFFER_ATOMIC_OR: 32b:tmp = MEM[ADDR];MEM[ADDR] |= DATA; RETURN_DATA = tmp.
 74 BUFFER_ATOMIC_XOR: 32b:tmp = MEM[ADDR];MEM[ADDR] ^= DATA; RETURN_DATA = tmp.
 75 BUFFER_ATOMIC_INC: 32b:tmp = MEM[ADDR];MEM[ADDR] = (tmp >= DATA) ? 0 : tmp + 1 (unsigned compare); RETURN_DATA = tmp.
 76 BUFFER_ATOMIC_DEC: 32b:tmp = MEM[ADDR];MEM[ADDR] = (tmp == 0 || tmp > DATA) ? DATA : tmp - 1 (unsigned compare); RETURN_DATA = tmp
 77 – 95 reserved.
-

80 – 95 Atomic (double).

- 96 BUFFER_ATOMIC_SWAP_X2: 64b:tmp = MEM[ADDR];MEM[ADDR] = DATA[0:1]; RETURN_DATA[0:1] = tmp.
 97 BUFFER_ATOMIC_CMPSWAP_X2: 64b:tmp = MEM[ADDR];src = DATA[0:1]; cmp = DATA[2:3]; MEM[ADDR] = (tmp == cmp) ? src : tmp; RETURN_DATA[0:1] = tmp.
-

Untyped Vector Memory Buffer Operation

98	BUFFER_ATOMIC_ADD_X2: 64b:tmp = MEM[ADDR];MEM[ADDR] += DATA[0:1]; RETURN_DATA[0:1] = tmp.
99	BUFFER_ATOMIC_SUB_X2: 64b:tmp = MEM[ADDR];MEM[ADDR] -= DATA[0:1]; RETURN_DATA[0:1] = tmp.
100	BUFFER_ATOMIC_SMIN_X2: 64b:tmp = MEM[ADDR];MEM[ADDR] -= (DATA[0:1] < tmp) ? DATA[0:1] : tmp (signed compare); RETURN_DATA[0:1] = tmp.
101	BUFFER_ATOMIC_UMIN_X2: 64b:tmp = MEM[ADDR];MEM[ADDR] -= (DATA[0:1] < tmp) ? DATA[0:1] : tmp (unsigned compare); RETURN_DATA[0:1] = tmp.
102	BUFFER_ATOMIC_SMAX_X2: 64b:tmp = MEM[ADDR];MEM[ADDR] -= (DATA[0:1] > tmp) ? DATA[0:1] : tmp (signed compare); RETURN_DATA[0:1] = tmp.
103	BUFFER_ATOMIC_UMAX_X2: 64b:tmp = MEM[ADDR];MEM[ADDR] -= (DATA[0:1] > tmp) ? DATA[0:1] : tmp (unsigned compare); RETURN_DATA[0:1] = tmp.
104	BUFFER_ATOMIC_AND_X2: 64b:tmp = MEM[ADDR];MEM[ADDR] &= DATA[0:1]; RETURN_DATA[0:1] = tmp.
105	BUFFER_ATOMIC_OR_X2: 64b:tmp = MEM[ADDR];MEM[ADDR] = DATA[0:1]; RETURN_DATA[0:1] = tmp.
106	BUFFER_ATOMIC_XOR_X2: 64b:tmp = MEM[ADDR];MEM[ADDR] ^= DATA[0:1]; RETURN_DATA[0:1] = tmp.
107	BUFFER_ATOMIC_INC_X2: 64b:tmp = MEM[ADDR];MEM[ADDR] = (tmp >= DATA[0:1]) ? 0 : tmp + 1 (unsigned compare); RETURN_DATA[0:1] = tmp.
108	BUFFER_ATOMIC_DEC_X2: 64b:tmp = MEM[ADDR];MEM[ADDR] = (tmp == 0 tmp > DATA[0:1]) ? DATA[0:1] : tmp - 1 (unsigned compare); RETURN_DATA[0:1] = tmp.

All other values are reserved.

reserved	25	
		Reserved.
ENCODING	[31:26]	enum(6) Must be 1 1 1 0 0 0.
VADDR	[39:32]	enum(8) VGPR address source. Can carry an offset or an index or both (can read two VGPRs).
VDATA	[47:40]	enum(8) Vector GPR to read/write result to.
SRSRC	[52:48]	enum(5) Scalar GPR that specifies the resource constant, in units of four SGPRs.
reserved	[54:53]	
		Reserved.
TFE	55	enum(1) Texture Fail Enable (for partially resident textures).

Untyped Vector Memory Buffer Operation

SOFFSET	[63:56]	enum(6)
	Byte offset added to the memory address. Scalar or constant GPR containing the base offset. This is always sent.	
	0 – 101 SGPR0 to SGPR101: Scalar general-purpose registers.	
	102 FLAT_SCRATCH_LO.	
	103 FLAT_SCRATCH_HI.	
	104 XNACK_MASK_LO. Carrizo APU only.	
	105 XNACK_MASK_HI. Carrizo APU only.	
	106 VCC_LO: vcc[31:0].	
	107 VCC_HI: vcc[63:32].	
	108 TBA_LO: Trap handler base address [31:0].	
	109 TBA_HI: Trap handler base address [63:32].	
	110 TMA_LO: Pointer to data in memory used by trap handler.	
	111 TMA_HI: Pointer to data in memory used by trap handler.	
	112 - 123 TTMP0 to TTMP11: Trap handler temporary registers (privileged).	
	124 M0. Memory register 0.	
	125 reserved.	
	126 EXEC_LO: exec[31:0].	
	127 EXEC_HI: exec[63:32].	
	128 0.	
	129 – 192: Signed integer 1 to 64.	
	193 – 208: Signed integer -1 to -16.	
	209 – 250: reserved.	

Typed Memory Buffer Operation

<i>Format</i>	MTBUF		
<i>Description</i>	Typed memory buffer operation. Two words		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	OFFSET	[11:0]	int(12)
		Unsigned byte offset.	
	OFFEN	12	enum(1)
		If set, send VADDR as an offset. If clear, use zero instead of an offset from a VGPR.	
	IDXEN	13	enum(1)
		If set, send VADDR as an index. If clear, treat the index as zero.	
	GLC	14	enum(1)
		If set, operation is globally coherent.	
	OP	[18:15]	enum(4)
	0	TBUFFER_LOAD_FORMAT_X: Untyped buffer load 1 Dword with format conversion.	
	1	TBUFFER_LOAD_FORMAT_XY: Untyped buffer load 2 Dwords with format conversion.	
	2	TBUFFER_LOAD_FORMAT_XYZ: Untyped buffer load 3 Dwords with format conversion.	
	3	TBUFFER_LOAD_FORMAT_XYZW: Untyped buffer load 4 Dwords with format conversion.	
	4	TBUFFER_STORE_FORMAT_X: Untyped buffer store 1 Dword with format conversion.	
	5	TBUFFER_STORE_FORMAT_XY: Untyped buffer store 2 Dwords with format conversion.	
	6	TBUFFER_STORE_FORMAT_XYZ: Untyped buffer store 3 Dwords with format conversion.	
	7	TBUFFER_STORE_FORMAT_XYZW: Untyped buffer store 4 Dwords with format conversion.	
	8	TBUFFER_LOAD_FORMAT_D16_X: Typed buffer load 1 dword with format conversion.	
	9	TBUFFER_LOAD_FORMAT_D16_XY: Typed buffer load 2 dwords with format conversion.	
	10	TBUFFER_LOAD_FORMAT_D16_XYZ: Typed buffer load 3 dwords with format conversion.	
	11	TBUFFER_LOAD_FORMAT_D16_XYZW: Typed buffer load 4 dwords with format conversion.	
	12	TBUFFER_STORE_FORMAT_D16_X: Typed buffer store 1 dword with format conversion.	
	13	TBUFFER_STORE_FORMAT_D16_XY: Typed buffer store 2 dwords with format conversion.	
	14	TBUFFER_STORE_FORMAT_D16_XYZ: Typed buffer store 3 dwords with format conversion.	
	15	TBUFFER_STORE_FORMAT_D16_XYZW: Typed buffer store 4 dwords with format conversion.	

Typed Memory Buffer Operation

DFMT	[22:19]	enum(4)
	Data format for typed buffer.	
	0 invalid	8 10_10_10_2
	1 8	9 2_10_10_10
	2 16	10 8_8_8_8
	3 8_8	11 32_32
	4 32	12 16_16_16_16
	5 16_16	13 32_32_32
	6 10_11_11	14 32_32_32_32
NFMT	[25:23]	enum(3)
	Number format for typed buffer.	
	0 unorm	
	1 snorm	
	2 uscaled	
	3 sscaled	
	4 uint	
	5 sint	
	6 reserved	
	7 float	
Encoding	[31:26]	enum(7)
	Must be 1 1 1 0 1 0.	
VADDR	[39:32]	enum(8)
	VGPR address source. Can carry an offset or an index or both (can read two successive VGPRs).	
VDATA	[47:40]	enum(8)
	Vector GPR to read/write result to.	
SRSRC	[52:48]	enum(5)
	Scalar GPR that specifies the resource constant, in units of four SGPRs.	
reserved	53	
	Reserved.	
SLC	54	enum(1)
	System Level Coherent.	
TFE	55	enum(1)

Typed Memory Buffer Operation

Texture Fail Enable (for partially resident textures).	
SOFFSET	[63:56] enum(6)
	Byte offset added to the memory address. Scalar or constant GPR containing the base offset. This is always sent.
	0 – 101 SGPR0 to SGPR101: Scalar general-purpose registers.
	102 FLAT_SCRATCH_LO.
	103 FLAT_SCRATCH_HI.
	104 XNACK_MASK_LO. Carrizo APU only.
	105 XNACK_MASK_HI. Carrizo APU only.
	106 VCC_LO: vcc[31:0].
	107 VCC_HI: vcc[63:32].
	108 TBA_LO: Trap handler base address [31:0].
	109 TBA_HI: Trap handler base address [63:32].
	110 TMA_LO: Pointer to data in memory used by trap handler.
	111 TMA_HI: Pointer to data in memory used by trap handler.
	112 – 123 TIMP0 to TIMP11: Trap handler temporary registers (privileged).
	124 MO. Memory register 0.
	125 reserved.
	126 - 127 reserved.
	128 0.
	129 – 192: Signed integer 1 to 64.
	193 – 208: Signed integer -1 to -16.
	209 – 255 reserved.

13.7 Vector Memory Image Instruction

Image Memory Buffer Operations

<i>Format</i>	MIMG		
<i>Description</i>	Image memory buffer operations. Two words.		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	reserved	[7:0]	Reserved.
	DMASK	[11:8]	enum(4) Enable mask for image read/write data components. bit0 = red, 1 = green, 2 = blue, 3 = alpha. At least one bit must be on. Data is assumed to be packed into consecutive VGPRs.
	UNORM	12	enum(1) When set to 1, forces the address to be un-normalized, regardless of T#. Must be set to 1 for image stores and atomics
	GLC	13	enum(1) If set, operation is globally coherent.
	DA	14	enum(1) Declare an Array. 1 Kernel has declared this resource to be an array of texture maps. 0 Kernel has declared this resource to be a single texture map.
	R128	15	enum(1) Texture resource size: 1 = 128b, 0 = 256b.
	TFE	16	enum(1) Texture Fail Enable (for partially resident textures).
	LWE	17	enum(1) LOD Warning Enable (for partially resident textures).
	OP	[24:18]	enum(8) 0 IMAGE_LOAD: Image memory load with format conversion specified in T#. No sampler. 1 IMAGE_LOAD_MIP: Image memory load with user-supplied mip level. No sampler. 2 IMAGE_LOAD_PCK: Image memory load with no format conversion. No sampler. 3 IMAGE_LOAD_PCK_SGN: Image memory load with no format conversion and sign extension. No sampler. 4 IMAGE_LOAD_MIP_PCK: Image memory load with user-supplied mip level, no format conversion. No sampler. 5 IMAGE_LOAD_MIP_PCK_SGN: Image memory load with user-supplied mip level, no format conversion and with sign extension. No sampler. 6 – 7 reserved. 8 IMAGE_STORE: Image memory store with format conversion specified in T#. No sampler. 9 IMAGE_STORE_MIP: Image memory store with format conversion specified in T# to user specified mip level. No sampler. 10 IMAGE_STORE_PCK: Image memory store of packed data without format conversion. No sampler.

Image Memory Buffer Operations

- 11 `IMAGE_STORE_MIP_PCK`: Image memory store of packed data without format conversion to user-supplied mip level. No sampler.
 - 12 – 13 reserved.
 - 14 `IMAGE_GET_RESINFO`: return resource info. No sampler.
 - 15 `IMAGE_ATOMIC_SWAP`: `dst=src`, returns previous value if `glc==1`.
 - 16 `IMAGE_ATOMIC_CMPSWAP`: `dst = (dst==cmp) ? src : dst`. Returns previous value if `glc==1`.
 - 17 `IMAGE_ATOMIC_ADD`: `dst += src`. Returns previous value if `glc==1`.
 - 18 `IMAGE_ATOMIC_SUB`: `dst -= src`. Returns previous value if `glc==1`.
 - 19 reserved.
 - 20 `IMAGE_ATOMIC_SMIN`: `dst = (src < dst) ? src : dst` (signed). Returns previous value if `glc==1`.
 - 21 `IMAGE_ATOMIC_UMIN`: `dst = (src < dst) ? src : dst` (unsigned). Returns previous value if `glc==1`.
 - 22 `IMAGE_ATOMIC_SMAX`: `dst = (src > dst) ? src : dst` (signed). Returns previous value if `glc==1`.
 - 23 `IMAGE_ATOMIC_UMAX`: `dst = (src > dst) ? src : dst` (unsigned). Returns previous value if `glc==1`.
 - 24 `IMAGE_ATOMIC_AND`: `dst &= src`. Returns previous value if `glc==1`.
 - 25 `IMAGE_ATOMIC_OR`: `dst |= src`. Returns previous value if `glc==1`.
 - 26 `IMAGE_ATOMIC_XOR`: `dst ^= src`. Returns previous value if `glc==1`.
 - 27 `IMAGE_ATOMIC_INC`: `dst = (dst >= src) ? 0 : dst+1`. Returns previous value if `glc==1`.
 - 28 `IMAGE_ATOMIC_DEC`: `dst = ((dst==0 || (dst > src)) ? src : dst-1)`. Returns previous value if `glc==1`.
 - 29 - 31 reserved.
 - 32 `IMAGE_SAMPLE`: sample texture map.
 - 33 `IMAGE_SAMPLE_CL`: sample texture map, with LOD clamp specified in shader.
 - 34 `IMAGE_SAMPLE_D`: sample texture map, with user derivatives.
 - 35 `IMAGE_SAMPLE_D_CL`: sample texture map, with LOD clamp specified in shader, with user derivatives.
 - 36 `IMAGE_SAMPLE_L`: sample texture map, with user LOD.
 - 37 `IMAGE_SAMPLE_B`: sample texture map, with lod bias.
 - 38 `IMAGE_SAMPLE_B_CL`: sample texture map, with LOD clamp specified in shader, with lod bias.
 - 39 `IMAGE_SAMPLE_LZ`: sample texture map, from level 0.
 - 40 `IMAGE_SAMPLE_C`: sample texture map, with PCF.
 - 41 `IMAGE_SAMPLE_C_CL`: `SAMPLE_C`, with LOD clamp specified in shader.
 - 42 `IMAGE_SAMPLE_C_D`: `SAMPLE_C`, with user derivatives.
 - 43 `IMAGE_SAMPLE_C_D_CL`: `SAMPLE_C`, with LOD clamp specified in shader, with user derivatives.
 - 44 `IMAGE_SAMPLE_C_L`: `SAMPLE_C`, with user LOD.
 - 45 `IMAGE_SAMPLE_C_B`: `SAMPLE_C`, with lod bias.
 - 46 `IMAGE_SAMPLE_C_B_CL`: `SAMPLE_C`, with LOD clamp specified in shader, with lod bias.
-

Image Memory Buffer Operations

- 47 IMAGE_SAMPLE_C_LZ: SAMPLE_C, from level 0.
 - 48 IMAGE_SAMPLE_O: sample texture map, with user offsets.
 - 49 IMAGE_SAMPLE_CL_O: SAMPLE_O with LOD clamp specified in shader.
 - 50 IMAGE_SAMPLE_D_O: SAMPLE_O, with user derivatives.
 - 51 IMAGE_SAMPLE_D_CL_O: SAMPLE_O, with LOD clamp specified in shader, with user derivatives.
 - 52 IMAGE_SAMPLE_L_O: SAMPLE_O, with user LOD.
 - 53 IMAGE_SAMPLE_B_O: SAMPLE_O, with lod bias.
 - 54 IMAGE_SAMPLE_B_CL_O: SAMPLE_O, with LOD clamp specified in shader, with lod bias.
 - 55 IMAGE_SAMPLE_LZ_O: SAMPLE_O, from level 0.
 - 56 IMAGE_SAMPLE_C_O: SAMPLE_C with user specified offsets.
 - 57 IMAGE_SAMPLE_C_CL_O: SAMPLE_C_O, with LOD clamp specified in shader.
 - 58 IMAGE_SAMPLE_C_D_O: SAMPLE_C_O, with user derivatives.
 - 59 IMAGE_SAMPLE_C_D_CL_O: SAMPLE_C_O, with LOD clamp specified in shader, with user derivatives.
 - 60 IMAGE_SAMPLE_C_L_O: SAMPLE_C_O, with user LOD.
 - 61 IMAGE_SAMPLE_C_B_O: SAMPLE_C_O, with lod bias.
 - 62 IMAGE_SAMPLE_C_B_CL_O: SAMPLE_C_O, with LOD clamp specified in shader, with lod bias.
 - 63 IMAGE_SAMPLE_C_LZ_O: SAMPLE_C_O, from level 0.
 - 64 IMAGE_GATHER4: gather 4 single component elements (2x2).
 - 65 IMAGE_GATHER4_CL: gather 4 single component elements (2x2) with user LOD clamp.
 - 66 IMAGE_GATHER4_L: gather 4 single component elements (2x2) with user LOD.
 - 67 IMAGE_GATHER4_B: gather 4 single component elements (2x2) with user bias.
 - 68 IMAGE_GATHER4_B_CL: gather 4 single component elements (2x2) with user bias and clamp.
 - 69 IMAGE_GATHER4_LZ: gather 4 single component elements (2x2) at level 0.
 - 70 IMAGE_GATHER4_C: gather 4 single component elements (2x2) with PCF.
 - 71 IMAGE_GATHER4_C_CL: gather 4 single component elements (2x2) with user LOD clamp and PCF.
 - 72 – 75 reserved.
 - 76 IMAGE_GATHER4_C_L: gather 4 single component elements (2x2) with user LOD and PCF.
 - 77 IMAGE_GATHER4_C_B: gather 4 single component elements (2x2) with user bias and PCF.
 - 78 IMAGE_GATHER4_C_B_CL: gather 4 single component elements (2x2) with user bias, clamp and PCF.
 - 79 IMAGE_GATHER4_C_LZ: gather 4 single component elements (2x2) at level 0, with PCF.
 - 80 IMAGE_GATHER4_O: GATHER4, with user offsets.
 - 81 IMAGE_GATHER4_CL_O: GATHER4_CL, with user offsets.
 - 82 – 83 reserved.
 - 84 IMAGE_GATHER4_L_O: GATHER4_L, with user offsets.
 - 85 IMAGE_GATHER4_B_O: GATHER4_B, with user offsets.
 - 86 IMAGE_GATHER4_B_CL_O: GATHER4_B_CL, with user offsets.
 - 87 IMAGE_GATHER4_LZ_O: GATHER4_LZ, with user offsets.
-

Image Memory Buffer Operations

- 88 IMAGE_GATHER4_C_O: GATHER4_C, with user offsets.
 89 IMAGE_GATHER4_C_CL_O: GATHER4_C_CL, with user offsets.
 90 – 91 reserved.
 92 IMAGE_GATHER4_C_L_O: GATHER4_C_L, with user offsets.
 93 IMAGE_GATHER4_C_B_O: GATHER4_B, with user offsets.
 94 IMAGE_GATHER4_C_B_CL_O: GATHER4_B_CL, with user offsets.
 95 IMAGE_GATHER4_C_LZ_O: GATHER4_C_LZ, with user offsets.
 96 IMAGE_GET_LOD: Return calculated LOD.
 97 – 103 reserved.
 104 IMAGE_SAMPLE_CD: sample texture map, with user derivatives (LOD per quad)
 105 IMAGE_SAMPLE_CD_CL: sample texture map, with LOD clamp specified in
 shader, with user derivatives (LOD per quad).
 106 IMAGE_SAMPLE_C_CD: SAMPLE_C, with user derivatives (LOD per quad).
 107 IMAGE_SAMPLE_C_CD_CL: SAMPLE_C, with LOD clamp specified in shader,
 with user derivatives (LOD per quad).
 108 IMAGE_SAMPLE_CD_O: SAMPLE_O, with user derivatives (LOD per quad).
 109 IMAGE_SAMPLE_CD_CL_O: SAMPLE_O, with LOD clamp specified in shader,
 with user derivatives (LOD per quad).
 110 IMAGE_SAMPLE_C_CD_O: SAMPLE_C_O, with user derivatives (LOD per quad).
 111 IMAGE_SAMPLE_C_CD_CL_O: SAMPLE_C_O, with LOD clamp specified in
 shader, with user derivatives (LOD per quad).
 All other values are reserved.

SLC	25	enum(1)	System Level Coherent.
ENCODING	[31:26]	enum(7)	Must be 1 1 1 1 0 0.
VADDR	[39:32]	enum(8)	Address source. Can carry an offset or an index. Specifies the VGPR that holds the first of the image address values.
VDATA	[47:40]	enum(8)	Vector GPR to which the result is written.
SRSRC	[52:48]	enum(5)	Scalar GPR that specifies the resource constant, in units of four SGPRs.
SSAMP	[57:53]	enum(5)	Scalar GPR that specifies the sampler constant, in units of four SGPRs.
reserved	[62:58]		Reserved.
D16	63	enum(1)	Convert 32-bit data to 16-bit data.

13.8 Export Instruction

Export

<i>Format</i>	EXP		
<i>Description</i>	Export (output) pixel color, pixel depth, vertex position, or vertex parameter data. Two words.		
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i>	<i>Format</i>
	EN	[3:0]	int(4)
	<p>This bitmask determines which VSRC registers export data.</p> <p>When COMPR is 0: VSRC0 only exports data when en[0] is set to 1; VSRC1 when en[1], VSRC2 when en[2], and VSRC3 when en[3].</p> <p>When COMPR is 1: VSRC0 contains two 16-bit data and only exports when en[0] is set to 1; VSRC1 only exports when en[2] is set to 1; en[1] and en[3] are ignored when COMPR is 1.</p>		
	TGT	[9:4]	enum(6)
	<p>Export target based on the enumeration below.</p> <p>0–7 EXP_MRT = Output to color MRT 0. Increment from here for additional MRTs. There are EXP_NUM_MRT MRTs in total.</p> <p>8 EXP_MRTZ = Output to Z.</p> <p>9 EXP_NULL = Output to NULL.</p> <p>12–15 EXP_POS = Output to position 0. Increment from here for additional positions. There are EXP_NUM_POS positions in total.</p> <p>32–63 EXP_PARAM = Output to parameter 0. Increment from here for additional parameters. There are EXP_NUM_PARAM parameters in total.</p> <p>All other values are reserved.</p>		
	COMPR	10	enum(1)
	Boolean. If true, data is exported in float16 format; if false, data is 32 bit.		
	DONE	11	enum(1)
	If set, this is the last export of a given type. If this is set for a color export (PS only), then the valid mask must be present in the EXEC register.		
	VM	12	enum(1)
	Mask contains valid-mask when set; otherwise, mask is just write-mask. Used only for pixel(mrt) exports.		
	reserved	[25:13]	
	Reserved.		
	ENCODING	[31:26]	enum(7)
	Must be 1 1 0 0 0 1.		
	VSRC0	[39:32]	enum(8)
	VGPR of the first data to export.		
	VSRC1	[47:40]	enum(8)
	VGPR of the second data to export.		
	VSRC2	[55:48]	enum(8)
	VGPR of the third data to export.		
	VSRC3	[63:56]	enum(8)
	VGPR of the fourth data to export.		

13.9 FLAT Instruction

Flat		
<i>Format</i>	FLAT	
<i>Description</i>	Export (output) pixel color, pixel depth, vertex position, or vertex parameter data. Two words.	
<i>Opcode</i>	<i>Field Name</i>	<i>Bits</i> <i>Format</i>
	reserved	[15:0] Reserved
	GLC	16 enum(1) If set, operation is globally coherent.
	SLC	17 enum(1) System Level Coherent. When set, indicates that the operation is "system level coherent". This controls the L2 cache policy.
	OP	[24:18] enum(7) 0 - 7 reserved. 8 FLAT_LOAD_UBYTE: Flat load unsigned byte. Zero extend to VGPR destination. 9 FLAT_LOAD_SBYTE: Flat load signed byte. Sign extend to VGPR destination. 10 FLAT_LOAD_USHORT: Flat load unsigned short. Zero extend to VGPR destination. 11 FLAT_LOAD_SSHORT: Flat load signed short. Sign extend to VGPR destination. 12 FLAT_LOAD_DWORD: Flat load Dword. 13 FLAT_LOAD_DWORDX2: Flat load 2 Dwords. 14 FLAT_LOAD_DWORDX4: Flat load 4 Dwords. 15 FLAT_LOAD_DWORDX3: Flat load 3 Dwords. 16 - 23 reserved. 24 FLAT_STORE_BYTE: Flat store byte. 25 reserved. 26 FLAT_STORE_SHORT: Flat store short. 27 reserved. 28 FLAT_STORE_DWORD: Flat store Dword. 29 FLAT_STORE_DWORDX2: Flat store 2 Dwords. 30 FLAT_STORE_DWORDX4: Flat store 4 Dwords. 31 FLAT_STORE_DWORDX3: Flat store 3 Dwords. 32 - 47 reserved. 48 FLAT_ATOMIC_SWAP: 32b. dst=src, returns previous value if rtn==1. 49 FLAT_ATOMIC_CMPSWAP: 32b, dst = (dst==cmp) ? src : dst. Returns previous value if rtn==1. src comes from the first data-VGPR, cmp from the second. 50 FLAT_ATOMIC_ADD: 32b, dst += src. Returns previous value if rtn==1. 51 FLAT_ATOMIC_SUB: 32b, dst -= src. Returns previous value if rtn==1. 52 reserved. 53 FLAT_ATOMIC_SMIN: 32b, dst = (src < dst) ? src : dst (signed comparison). Returns previous value if rtn==1. 54 FLAT_ATOMIC_UMIN: 32b, dst = (src < dst) ? src : dst (unsigned comparison). Returns previous value if rtn==1. 55 FLAT_ATOMIC_SMAX: 32b, dst = (src > dst) ? src : dst (signed comparison). Returns previous value if rtn==1.

Flat

-
- 56 FLAT_ATOMIC_UMAX: 32b, dst = (src > dst) ? src : dst (unsigned comparison). Returns previous value if rtn==1.
- 57 FLAT_ATOMIC_AND: 32b, dst &= src. Returns previous value if rtn==1.
- 58 FLAT_ATOMIC_OR: 32b, dst |= src. Returns previous value if rtn==1.
- 59 FLAT_ATOMIC_XOR: 32b, dst ^= src. Returns previous value if rtn==1.
- 60 FLAT_ATOMIC_INC: 32b, dst = (dst >= src) ? 0 : dst+1 (unsigned comparison). Returns previous value if rtn==1.
- 61 FLAT_ATOMIC_DEC: 32b, dst = ((dst==0 || (dst > src)) ? src : dst-1 (unsigned comparison). Returns previous value if rtn==1.
- 62 – 79 reserved.
- 80 FLAT_ATOMIC_SWAP_X2: 64b, dst=src, returns previous value if rtn==1.
- 81 FLAT_ATOMIC_CMPSWAP_X2: 64b, dst = (dst==cmp) ? src : dst. Returns previous value if rtn==1. src comes from the first two data-VGPRs, cmp from the second two.
- 82 FLAT_ATOMIC_ADD_X2: 64b, dst += src. Returns previous value if rtn==1.
- 83 FLAT_ATOMIC_SUB_X2: 64b, dst -= src. Returns previous value if rtn==1.
- 84 reserved.
- 85 FLAT_ATOMIC_SMIN_X2: 64b, dst = (src < dst) ? src : dst (signed comparison). Returns previous value if rtn==1.
- 86 FLAT_ATOMIC_UMIN_X2: 64b, dst = (src < dst) ? src : dst (unsigned comparison). Returns previous value if rtn==1.
- 87 FLAT_ATOMIC_SMAX_X2: 64b, dst = (src > dst) ? src : dst (signed comparison). Returns previous value if rtn==1.
- 88 FLAT_ATOMIC_UMAX_X2: 64b, dst = (src > dst) ? src : dst (unsigned comparison). Returns previous value if rtn==1.
- 89 FLAT_ATOMIC_AND_X2: 64b, dst &= src. Returns previous value if rtn==1.
- 90 FLAT_ATOMIC_OR_X2: 64b, dst |= src. Returns previous value if rtn==1.
- 91 FLAT_ATOMIC_XOR_X2: 64b, dst ^= src. Returns previous value if rtn==1.
- 92 FLAT_ATOMIC_INC_X2: 64b, dst = (dst >= src) ? 0 : dst+1. Returns previous value if rtn==1.
- 93 FLAT_ATOMIC_DEC_X2: 64b, dst = ((dst==0 || (dst > src)) ? src : dst - 1. Returns previous value if rtn==1.

All other values are reserved.

ENCODING	[31:26]	enum(7)
	Must be 1 1 0 1 1 1.	
ADDR	[39:32]	enum(8)
	Source of flat address VGPR.	
DATA	[47:40]	enum(8)
	Source data.	
reserved	[54:48]	
	Reserved	
TFE	55	enum(1)
	Texture Fail Enable. For partially resident textures.	
VDST	[63:56]	enum(14)
	Destination VGPR.	
