



# Intel<sup>®</sup> OpenSource HD Graphics Programmer's Reference Manual (PRM) Volume 4 Part 2: Subsystem and Cores – Message Gateway, URB, and IS (SandyBridge)

For the 2011 Intel Core Processor Family

*May 2011*

*Revision 1.0*

**NOTICE:**

This document contains information on products in the design phase of development, and Intel reserves the right to add or remove product features at any time, with or without changes to this open source documentation.



## Creative Commons License

**You are free to Share** — to copy, distribute, display, and perform the work

**Under the following conditions:**

**Attribution.** You must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work).

**No Derivative Works.** You may not alter, transform, or build upon this work.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The SandyBridge chipset family, Havendale/Auburndale chipset family, Intel® 965 Express Chipset Family, Intel® G35 Express Chipset, and Intel® 965GMx Chipset Mobile Family Graphics Controller may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel® sales office or your distributor to obtain the latest specifications and before placing your product order. I2C is a two-wire communications bus/protocol developed by Philips. SMBus is a subset of the I2C bus/protocol and was developed by Intel®. Implementations of the I2C bus/protocol may require licenses from various entities, including Philips Electronics N.V. and North American Philips Corporation.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

\*Other names and brands may be claimed as the property of others.

**Copyright © 2011, Intel Corporation. All rights reserved.**



# Contents

<b>1. Message Gateway .....</b>	<b>7</b>
1.1 Messages.....	7
1.1.1 Message Descriptor .....	8
1.1.2 OpenGateway Message.....	9
1.1.3 CloseGateway Message .....	11
1.1.4 ForwardMsg Message.....	11
1.1.5 GetTimeStamp Message [DevSNB+] .....	14
1.1.6 BarrierMsg Message .....	15
<b>2. Unified Return Buffer (URB).....</b>	<b>18</b>
2.1 URB Size.....	18
2.2 URB Access .....	18
2.3 State.....	19
2.4 Messages [DevSNB].....	19
2.4.1 Execution Mask.....	20
2.4.2 Message Descriptor .....	20
2.4.3 URB_WRITE .....	23
2.4.4 FF_SYNC Messages .....	28
<b>3. Execution Unit ISA .....</b>	<b>31</b>
3.1 Introduction .....	31
3.1.1 Objective and Scope .....	31
3.1.2 Terms and Acronyms .....	31
3.1.3 Formats and Conventions.....	34
<b>4. EU Data Types .....</b>	<b>35</b>
4.1 Fundamental Data Types.....	35
4.2 Numerical Data Types.....	36
4.2.1 Unsigned Integers .....	38
4.2.2 Signed Integers .....	38
4.2.3 Single Precision Floating-Point Numbers .....	38
4.2.4 Packed Signed Half-Byte Integer Vector .....	38
4.2.5 Packed 8-bit Restricted Float Vector .....	39
4.3 Floating Point Modes .....	42
4.3.1 IEEE Floating Point Mode.....	42
4.3.2 Alternative Floating Point Mode .....	49
4.4 Type Conversion .....	50
4.4.1 Float to Integer .....	50
4.4.2 Integer to Integer with Same or Higher Precision .....	51
4.4.3 Integer to Integer with Lower Precision.....	51
4.4.4 Integer to Float .....	51
<b>5. Execution Environment .....</b>	<b>52</b>
5.1 Overview .....	52
5.2 Primary Usage Models.....	53
5.2.1 AOS and SOA Data Structures .....	53
5.2.2 SIMD4 Mode of Operation .....	54
5.2.3 SIMD4x2 Mode of Operation.....	55
5.2.4 SIMD16 Mode of Operation .....	57
5.2.5 SIMD8 Mode of Operation .....	58
5.3 Registers and Register Regions .....	59
5.3.1 Register Files .....	59
5.3.2 GRF Registers.....	59



5.3.3	MRF Registers .....	61
5.3.4	ARF Registers .....	61
5.3.5	Immediate.....	81
5.3.6	Region Parameters .....	82
5.3.7	Region Addressing Modes .....	88
5.3.8	Access Modes.....	93
5.3.9	Execution Data Type.....	94
5.3.10	Register Region Restrictions.....	94
5.3.11	Destination Operand Description .....	97
5.4	SIMD Execution Control.....	98
5.4.1	Predication .....	98
5.4.2	No Predication.....	99
5.4.3	Predication with Horizontal Combination .....	100
5.4.4	Predication with Vertical Combination .....	102
5.5	Instruction Compaction .....	102
5.5.1	Motivation and Expected Usage .....	102
5.5.2	Hardware Behavior .....	106
5.5.3	Rules and Restrictions .....	108
5.6	End of Thread .....	108
5.7	Creating Conditional Flags.....	109
5.8	Destination Hazard.....	111
5.9	Non-present Operands.....	112
5.10	Instruction Prefetch .....	112
<b>6.</b>	<b>Exceptions .....</b>	<b>113</b>
6.1	Introduction .....	113
6.2	Exception-Related Architectural Registers .....	113
6.3	System Routine.....	114
6.3.1	General Flow of the System Routine .....	114
6.3.2	Invoking the System Routine .....	115
6.3.3	Returning to the Application Thread .....	115
6.3.4	System-IP (SIP).....	116
6.3.5	System Routine Register Space .....	116
6.3.6	System-Scratch Memory Space.....	117
6.3.7	Conditional Instructions Within System Routines .....	117
6.3.8	Messages in System Routines.....	118
6.3.9	Use of 'NoDDClr' .....	118
6.4	Exception Descriptions.....	119
6.4.1	'Illegal' opcode.....	119
6.4.2	Undefined opcode .....	119
6.4.3	Software Exception .....	119
6.4.4	Breakpoint .....	119
6.4.5	External Halt.....	119
6.5	Events Which Do Not Generate Exceptions .....	120
6.6	System Handler Example.....	121
<b>7.</b>	<b>Instruction Set Summary.....</b>	<b>127</b>
7.1	Instruction Set Characteristics .....	127
7.1.1	SIMD Instructions and SIMD Width.....	127
7.1.2	Instruction Operands and Register Regions .....	127
7.1.3	Instruction Execution.....	128
7.2	Instruction Machine Formats.....	128
7.2.1	Common Instruction Fields .....	132
7.2.2	Instruction Operation Doubleword (DW0).....	139
7.2.3	Instruction Destination Doubleword (DW1).....	144



7.2.4	Instruction Source-0 Doubleword (DW2)	148
7.2.5	Instruction Source-1 Doubleword (DW3)	153
7.3	Opcode Encoding	156
7.3.1	Move and Logic Instructions	156
7.3.2	Flow Control Instructions	158
7.3.3	Miscellaneous Instructions	158
7.3.4	Parallel Arithmetic Instructions	159
7.3.5	Vector Arithmetic Instructions	160
7.3.6	Special Instructions	161
7.4	Native Instruction BNF	162
7.4.1	Instruction Groups	162
7.4.2	Destination Register	163
7.4.3	Source Register	164
7.4.4	Address Registers	165
7.4.5	Register Files and Register Numbers	166
7.4.6	Relative Location and Stack Control	167
7.4.7	Regions	167
7.4.8	Types	168
7.4.9	Write Mask	168
7.4.10	Swizzle Control	168
7.4.11	Immediate Values	168
7.4.12	Predication and Modifiers	169
7.4.13	Instruction Options	170
7.5	Deprecated Features	171
7.5.1	Defeatured Instructions	171
7.5.2	Others	171
<b>8.</b>	<b>Instruction Set Reference</b>	<b>172</b>
8.1	Conventions	172
8.1.1	Pseudo Code Format	172
8.1.2	General Macros and Definitions	172
8.2	Evaluate Write Enable	173
8.3	Instruction Description	174
8.3.1	add – Addition	174
8.3.2	and – Logical And	177
8.3.3	asr – Arithmetic Shift Right	179
8.3.4	avg – Average	180
8.3.5	break – Break [DevSNB]	181
8.3.6	call – Call [DevSNB]	182
8.3.7	case – Case [DevSNB]	184
8.3.8	cmp – Compare	187
8.3.9	cmpn – Compare NaN	188
8.3.10	cont – Continue [DevSNB]	190
8.3.11	dp2 – Dot Product 2	191
8.3.12	dp3 – Dot Product 3	192
8.3.13	dp4 – Dot Product 4	194
8.3.14	dph –Dot Product Homogeneous	195
8.3.15	else – Else [DevSNB]	196
8.3.16	endif – End-If	197
8.3.17	fork – Fork [DevSNB]	199
8.3.18	frc – Fraction	201
8.3.19	halt – Halt [DevSNB]	203
8.3.20	if – If [DevSNB]	204
8.3.21	jmpI – Jump Indexed	206



8.3.22	line – Line .....	208
8.3.23	lzd – Leading Zero Detection .....	209
8.3.24	lrp – Linear Interpolation .....	210
8.3.25	mac – Multiply Accumulate .....	211
8.3.26	mach – Multiply Accumulate High .....	212
8.3.27	mad – Multiply Add .....	214
8.3.28	math – Extended Math Function [DevSNB] .....	215
8.3.29	mov – Move .....	218
8.3.30	movi – Move Indexed .....	219
8.3.31	mul – Multiply .....	221
8.3.32	nop – No Operation .....	224
8.3.33	not – Logic Not .....	225
8.3.34	or – Logic Or .....	226
8.3.35	pln – Plane .....	227
8.3.36	ret – Return [DevSNB] .....	229
8.3.37	rndd – Round Down .....	230
8.3.38	rndu – Round Up .....	232
8.3.39	rnde – Round to Even .....	234
8.3.40	rndz – Round to Zero .....	236
8.3.41	sad2 – Sum of Absolute Difference 2 .....	238
8.3.42	sada2 – Sum of Absolute Difference Accumulate 2 .....	239
8.3.43	sel – Select .....	240
8.3.44	send – Send Message [DevSNB] .....	242
8.3.45	sendc – Conditional Send Message .....	247
8.3.46	shl – Shift Left .....	248
8.3.47	shr – Shift Right .....	249
8.3.48	wait – Wait Notification .....	250
8.3.49	while – While [DevSNB] .....	251
8.3.50	xor – Logic Xor .....	253
<b>9.</b>	<b>EU Programming Guide .....</b>	<b>255</b>
9.1	Assembler Pragmas .....	255
9.1.1	Declarations .....	255
9.1.2	Defaults and Defines .....	256
9.1.3	Example Pragma Usages .....	257
9.1.4	Assembly Programming Guideline .....	260
9.2	Usage Examples .....	262
9.2.1	Vector Immediate .....	262
9.2.2	Destination Mask for DP4 and Destination Dependency Control .....	263
9.2.3	Null Register as the Destination .....	263
9.2.4	Use of LINE Instruction .....	264
9.2.5	Mask for SEND Instruction .....	265
9.2.6	Flow Control Instructions .....	268
9.2.7	Execution Masking .....	269



# 1. Message Gateway

The Message Gateway shared function provides a mechanism for active thread-to-thread communication. Such thread-to-thread communication is based on direct register access. One thread, a **requester thread**, is capable of writing into the GRF register space of another thread, a **recipient thread**. Such direct register access between two threads in a multi-processor environment some time is referred to as **remote register access**. Remote register access may include read or write. GEN4 architecture supports **remote register write**, but not remote register read (natively). Message Gateway facilitates such remote register write via message passing. The requester thread sends a message to Message Gateway requesting a write to the recipient thread's GRF register space. Message Gateway sends a writeback message to the recipient thread to complete the register write on behalf of the requester. The requester thread and the recipient thread may be on the same EU or on different EUs.

[DevSNB+] When Bypass Gateway Control is set to 1, commands **OpenGateway** and **CloseGateway** are no longer used, the gateway parameters are taking the default values as the following:

- **RegBase** = 0
- **Gateway Size** check and **Key** check are bypassed.
- **Gateway Open** (an internal signal that is used to be set by OpenGateway message) check is bypassed

## 1.1 Messages

Message Gateway supports such thread-to-thread communication with the following three messages:

- **OpenGateway**: opens a gateway for a requester thread. Once a thread successfully opens its gateway, it can be a recipient thread to receive remote register write.
- **CloseGateway**: closes the gateway for a requester thread. Once a thread successfully closes its gateway, Message Gateway will block any future remote register writes to this thread.
- **ForwardMsg**: forwards a formatted message (remote register write) from a requester thread to a recipient thread.
- **GetTimeStamp [DevSNB+]**: reads absolute and relative timestamps for a requester thread.
- **BarrierMsg [DevSNB+]**: a set of threads sends this message to the Gateway. When all threads in a group have sent the message, a reply (both a register write and an NO notification) is sent to each member of the group.
- **UpdateGatewayState [DevSNB+]**: updates the internal state of the Message Gateway. One example usage is to allow a control thread to change Barrier Byte to convey dynamic state information. This may be used to support interrupt (for advanced scheduler) when persistent compute/worker threads synchronized using Barrier.



### 1.1.1 Message Descriptor

The following message descriptor applies to all messages supported by Message Gateway.

Bit	Description
19	<b>Header Present</b> This bit must be set to <b>zero</b> for all Message Gateway messages.
18:17	Ignored
16:15	<b>Notify.</b> Send Notification Signal. This is a two-bit field indicating which notify event is sent. 00: No notify 01: Increment recipient thread's N0 notification counter 10: Increment recipient thread's N2 notification counter 11: Reserved  This field is only valid for a ForwardMsg message. It is ignored for other messages. The BarrierMsg message always increments the N0 notification counter.
14	<b>AckReq.</b> Acknowledgment Required. When this bit is set, an acknowledgment return message is required. Message Gateway will send a writeback message containing the error code to the requester thread using the post destination register address. When this bit is not set, no writeback message is sent to the requesting thread by Message Gateway, even if an error occurs.  This field is valid for OpenGateway, CloseGateway, ForwardMsg and BarrierMsg messages.  When this bit is set, post destination register must be valid and the response length must be 1.  When this bit is not set, post destination register must be <i>null</i> and the response length must be 0.  This bit cannot be set when EOT is set; otherwise, hardware behavior is undefined.  0 = No Acknowledgement is required. 1 = Acknowledgement is required.
13:3	Reserved: MBZ





Bit	Description
2:0	<p><b>SubFuncID.</b> Identify the supported sub-functions by Message Gateway. Encodings are:</p> <p>000 = <b>OpenGateway.</b> Open the gateway for the requester thread.</p> <p>001 = <b>CloseGateway.</b> Close the gateway for the requester thread.</p> <p>010 = <b>ForwardMsg.</b> Forward the formatted message to the recipient thread with the given offset from the recipient's register base.</p> <p>011 = <b>GetTimeStamp [DevSNB+].</b> Read absolute and relative timestamps.</p> <p>100 = <b>BarrierMsg.</b> Record an additional thread reaching the barrier.</p> <p>101 = <b>UpdateGatewayState.</b> Update the barrier byte for a barrier.</p> <p>110 = <b>MMIO Read/Write</b></p> <p>Others are reserved</p> <p>[Pre-DevSNB] 011 and 1xx = Reserved</p>

## 1.1.2 OpenGateway Message

The OpenGateway message opens a communication channel between the requesting thread and other threads. It specifies a key for other threads to access its gateway, as well as the GRF register range allowed to be written. The message consists of a single 256-bit message payload.

If the AckReq bit is set, a single 256-bit payload writeback message is sent back to the requesting thread after completion of the OpenGateway function. Only the least significant DWord in the post destination register is overwritten.

If the EOT is set for this message, Message Gateway will ignore this message; instead, it will close the gateway for the requesting thread regardless of the previous state of the gateway.

It is software's policy to determine how to generate the key.

### 1.1.2.1 Message Payload

DWord	Bit	Description
M0.7	31:0	<b>Reserved</b>
M0.6	31:0	<b>Reserved</b>
M0.5	31:29	Reserved: MBZ
	28:21	<p><b>RegBase:</b> The register base address to be stored in the Message Gateway. It is used to compute the destination GRF register address from the offset field in ForwardMsg. RegBase contains 256-bit GRF aligned register address.</p> <p>Note 1: This field aligns with bits [28:21] of the Offset field of the message payload for ForwardMsg.</p> <p>Note 2: the most significant bit of this field must be zero.</p> <p>Format = U8</p> <p>Range = [0,127]</p>



DWord	Bit	Description
	20:11	Reserved: MBZ
	10:8	[DevSNB+] This field is reserved: MBZ
	7:0	<b>Dispatch ID:</b> This ID is assigned by the fixed function unit and is a unique identifier for the thread. It is used to free up resources used by the thread upon thread completion. This field is ignored by Message Gateway This field is only required for a thread that is created by a fixed function (therefore, not a child thread) and EOT bit is set for the message.
M0.4	31:16	Reserved: MBZ
	15:0	[DevSNB+] This field is reserved: MBZ
M0.3:0		Ignored

### 1.1.2.2 Writeback Message to Requester Thread

The writeback message is only sent if the **AckReq** bit in the message descriptor is set.

DWord	Bit	Description
W0.7:1		Reserved (not overwritten)
W0.0	31:20	Reserved
	19:16	<b>Shared Function ID:</b> Contains the message gateway's shared function ID.
	15:3	Reserved
	2:0	<b>Error Code</b> 000: <i>Successful</i> . No Error (Normal) 001: <i>Gateway Size Exceeded</i> . 101: <i>Opcode Error</i> . Attempt to send a message which is not either open/close/forward other codes: Reserved



### 1.1.3 CloseGateway Message

The CloseGateway message closes a communication channel for the requesting thread that was previously opened with OpenGateway. Each thread is allowed to have only one open gateway at a time, thus no additional information in the message payload is required to close the gateway. The message consists of a single 256-bit message payload.

If the AckReq bit is set, a single 256-bit payload writeback message is sent back to the requesting thread after completion of the CloseGateway function. Only the least significant DWord in the post destination register is overwritten.

#### 1.1.3.1 Message Payload

DWord	Bit	Description
M0.7:6		Ignored
M0.5	31:8	Ignored
	7:0	<b>Dispatch ID:</b> This ID is assigned by the fixed function unit and is a unique identifier for the thread. It is used to free up resources used by the thread upon thread completion. This field is ignored by Message Gateway This field is only required for a thread that is created by a fixed function (therefore, not a child thread) and EOT bit is set for the message.
M0.4:0		Ignored

#### 1.1.3.2 Writeback Message to Requester Thread

The writeback message is only sent if the **AckReq** bit in the message descriptor is set.

DWord	Bit	Description
W0.7:1		Reserved (not overwritten)
W0.0	31:20	Reserved
	19:16	<b>Shared Function ID:</b> Contains the message gateway's shared function ID.
	15:3	Reserved
	2:0	<b>Error Code</b> 000: <i>Successful. No Error (Normal)</i> 101: <i>Opcode Error. Attempt to send a message which is not either open/close/forward</i> other codes: Reserved

### 1.1.4 ForwardMsg Message

The ForwardMsg message gives the ability for a requester thread to write a **data segment** in the form of a byte, a dword, 2 dwords, or 4 dwords to a GRF register in a recipient thread. The message consists of a single 256-bit message payload, which contains the specially formatted data segment.

The ForwardMsg message utilizes a communication channel previously opened by the recipient thread. The recipient thread has communicated its EUID, TID, and key to the requester thread previously via



some other mechanism. Generally, this is done through the thread spawn message from parent to child thread, allowing each child (requester) to then communicate with its parent through a gateway opened by the parent (recipient). The child could then use ForwardMsg message to communicate its own EUID, TID, and key back to the parent to enable bi-directional communication after opening its own gateway.

If the AckReq bit is set, a single 256-bit payload writeback message is sent back to the requester thread after completion of the ForwardMsg function. Only the least significant DWord in the post destination register is overwritten.

If the Notify bit in the message descriptor is set, a 'notification' is sent to the recipient thread in order to increment the recipient thread's notification counter. This allows multiple messages to be sent to the recipient without waking up the recipient thread. The last message, having this bit set, will then wake up the recipient thread.

### 1.1.4.1 Message Payload

DWord	Bit	Description
M0.7	31:0	<b>Reserved</b>
M0.6	31:0	<b>Reserved</b>
M0.5	31:29	Reserved: MBZ
	28:16	<b>Offset:</b> It provides the destination register position in the recipient thread GRF register space as the offset from the RegBase stored in the recipient thread's gateway entry. The offset is in unit of byte, such that bits [28:21] is the 256-bit aligned register offset and bits [4:0] is the sub-register offset. The sub-register offset must be aligned to the Length field in bits [10:8]. The subfields of Offset are further illustrated as the following. Offset[28:21]: Register offset from the gateway base (Range [0, 127]: bit 12 MBZ) Offset[20:18]: DW offset Offset[17:16]: Byte offset (must be 00 for all DW length cases) <b>Programming restriction:</b> R0 can not be used as destination GRF register for ForwardMsg. NULL register is also not allowed as destination.
	15:11	Reserved: MBZ
	10:8	<b>Length:</b> The length of the data segment. 000: 1 byte 001: 1 word 010: 1 dword 011: 2 dwords 100: 4 dwords 101-111: Reserved
	7:0	<b>Dispatch ID:</b> This ID is assigned by the fixed function unit and is a unique identifier for the thread. It is used to free up resources used by the thread upon thread completion. This field is ignored by Message Gateway This field is only required for a thread that is created by a fixed function (therefore, not a child thread) and EOT bit is set for the message.
M0.4	31:30	Ignored
	29	<b>Reserved</b>
	28	<b>Reserved.</b>



DWord	Bit	Description
	27:24	<b>EUID:</b> The Execution Unit ID as part of the Recipient field is used to identify the recipient thread to whom the message is forwarded.
	23:19	Ignored
	18:16	<b>TID:</b> The Thread ID as part of the Recipient field is used to identify the recipient thread to whom the message is forwarded.
	15:0	<b>Key</b> The key to match with the one stored in the recipient thread's entry in Message Gateway. [DevSNB+] Ignored
M0.3	31:0	<b>Data Segment DWord 3:</b> valid only for the 4-DWord data segment length
M0.2	31:0	<b>Data Segment DWord 2:</b> valid only for the 4-DWord data segment length
M0.1	31:0	<b>Data Segment Dword 1:</b> valid only for the 2- and 4-DWord data segment lengths
M0.0	31:24	<b>Data Segment Byte 0:</b> the same byte must be copied to all four positions within this DWord. Valid only for the 1-Byte data segment length.
	23:16	<b>Data Segment Byte 0</b>
	15:8	<b>Data Segment Byte 0</b>
	7:0	<b>Data Segment Byte 0</b>
		<b>Data Segment Dword 0:</b> valid only for the 1-, 2- and 4-Dword data segment lengths

### 1.1.4.2 Writeback Message to Requester Thread

The writeback message is only sent if the **AckReq** bit in the message descriptor is set.

DWord	Bit	Description
W0.7:1		Reserved (not overwritten)
W0.0	31:20	Reserved
	19:16	<b>Shared Function ID:</b> Contains the message gateway's shared function ID.
	15:3	Reserved
	2:0	<b>Error Code</b> 000: <i>Successful</i> . No Error (Normal) 001: Reserved 010: <i>Gateway Closed</i> . Attempt to send a message through a closed gateway 011: <i>Key Mismatched</i> . [Pre-DevSNB] Attempt to send a message with a mismatching key 100: <i>Limit Exceeded</i> . [Pre-DevSNB] Attempt to send a message with offset beyond the gateway limit 101: <i>Opcode Error</i> . Attempt to send a message which is not either open/close/forward 110: <i>Invalid Message Size</i> . Attempt to forward a message with length greater than 4 DW 111: Reserved



### 1.1.4.3 Writeback Message to Recipient Thread

This message contains the byte or dwords data segment indicated in the message written to the GRF register offset indicated. Only the byte/dword(s) will be enabled, all other data in the GRF register is untouched.

### 1.1.5 GetTimeStamp Message [DevSNB+]

The GetTimeStamp message gives the ability for a requester thread to read the timestamps back from the message gateway. The message consists of a single 256-bit message payload.

AbsoluteTimeLap is based on an absolute wall clock in unit of nSec/uSec that is independent of context switch or GPU frequency adjustment. Message Gateway shares the same GPU timestamp. Details can be found in the TIMESTAMP register section in *vol1c Memory Interface and Command Stream*.

RelativeTimeLap is based on a relative time count that is counting the GPU clocks for the context. The relative time count is saved/restored during context switch.

#### 1.1.5.1 Message Payload

DWord	Bit	Description
M0.7	31:0	Reserved
M0.6	31:0	Reserved
M0.5	31	<b>Return to High GRF:</b> 0: the return 128-bit data goes to the first half of the destination GRF register 1: the return 128-bit data goes to the second half of the destination GRF register
	30:8	Reserved : MBZ
	7:0	<b>Dispatch ID:</b> This ID is assigned by the fixed function unit and is a unique identifier for the thread. It is used to free up resources used by the thread upon thread completion. This field is ignored by Message Gateway This field is only required for a thread that is created by a fixed function (therefore, not a child thread) and EOT bit is set for the message.
M0.4	31:0	Ignored
M0.3	31:0	Ignored
M0.2	31:0	Ignored
M0.1	31:0	Ignored
M0.0	31:0	Ignored

#### 1.1.5.2 Writeback Message to Requester Thread

As the writeback message is only sent if the **AckReq** bit in the message descriptor is set, **AckReq** bit must be set for this message.

Only half of the destination GRF register is updated (via write-enables). The other half of the register is not changed. This is determined by the **Return to High GRF** control field.

Writeback Message if Return to High GRF is set to 0:



DWord	Bit	Description
W0.7:4		Reserved (not overwritten)
W0.3	31:0	<b>RelativeTimeLapHigh:</b> This field returns the MSBs of time lap for the relative clock since the previous reset. This field represents 1.024 uSec increment of the time stamp. Hardware handles the wraparound (over 64 bit boundary) of the timestamp. Format: U12
W0.2	31:20	<b>RelativeTimeLapLow:</b> This field returns the LSBs of time lap for the relative clock since the previous reset. This field represents 1/4 nSec increment of the time stamp. Hardware handles the wraparound (over 64 bit boundary) of the timestamp. Format: U12
	19:0	Reserved : MBZ
W0.1	31:0	<b>AbsoluteTimeLapHigh:</b> This field returns the MSBs of time lap for the absolute clock since the previous reset. This field represents 1.024 uSec increment of the time stamp. Hardware handles the wraparound (over 64 bit boundary) of the timestamp. Format: U12
W0.0	31:20	<b>AbsoluteTimeLapLow:</b> This field returns the LSBs of time lap for the absolute clock since the previous reset. This field represents 1/4 nSec increment of the time stamp. Hardware handles the wraparound (over 64 bit boundary) of the timestamp. Format: U12
	19:0	Reserved : MBZ

Writeback Message if Return to High GRF is set to 1:

DWord	Bit	Description
W0.7	31:0	<b>RelativeTimeLapHigh</b>
W0.6	31:20	<b>RelativeTimeLapLow</b>
	19:0	Reserved : MBZ
W0.5	31:0	<b>AbsoluteTimeLapHigh</b>
W0.4	31:20	<b>AbsoluteTimeLapLow</b>
	19:0	Reserved : MBZ
W0.3:0		Reserved : MBZ

### 1.1.6 BarrierMsg Message

The BarrierMsg message gives the ability for multiple threads to synchronize their progress. This is useful when there are data shared between threads. The message consists of a single 256-bit message payload.

Upon receiving one such message, Message Gateway increments the Barrier counter and mark the Barrier requester thread. There is no immediate response from the Message Gateway. When the counter value equates **Barrier Thread Count**, Message Gateway will send response back to all the Barrier requesters.



### 1.1.6.1 Message Payload

DWord	Bit	Description
M0.7	31:0	<b>Reserved</b>
M0.6	31:0	<b>Reserved</b>
M0.5	31:0	Ignored
M0.4	31:0	Ignored
M0.3	31:0	Ignored
M0.2	31:28	Ignored
	27:24	<b>BarrierID.</b> This field indicates which one from the 16 Barrier States is updated. Format: U4 Note: this field location matches with that of R0 header.
	23:16	<b>Barrier.Offset.</b> This is the offset for the Barrier to indicate the offset from the requester's RegBase. Regbase is 0 if Bypass Gateway Control is 1. Barrier.Offset + RegBase must be in the valid GRF range, but not point to r0. Otherwise, hardware behavior is undefined. It is in unit of 256-bit GRF register. The most significant bit of this field must be zero. Format = U8 Range = [0,127] [DevSNB] The least significant bit of this field is ignored by hardware, so only even offsets are allowed..
	15:0	Ignored
M0.1	31:0	Ignored
M0.0	31:4	Ignored

### 1.1.6.2 Writeback Message to Requester Thread

The writeback message is only sent if the **AckReq** bit in the message descriptor is set.

DWord	Bit	Description
W0.7:1		Reserved (not overwritten)
W0.0	31:20	Reserved
	19:16	<b>Shared Function ID:</b> Contains the message gateway's shared function ID.
	15:3	Reserved
	2:0	<b>Error Code</b> 000: <i>Successful.</i> No Error (Normal) 001: <i>Error (Barrier is inactive).</i> Other encodings are reserved





### 1.1.6.3 Broadcast Writeback Message

[DevSNB] When the count for a Barrier (identified by BarrierID) reaches Barrier.Count, Message Gateway broadcasts the following messages to all threads using the BarrierID.

This message contains one single byte written to the GRF register at the *Barrier.Offset + RegBase* location. Only the byte will be enabled, all other data in the GRF register is untouched. In addition, the message will have the notification bit set for N0 so that the kernel set a dependency on either the register write or N0.

DWord	Bit	Description
W0.7:1		Reserved (not overwritten)
W0.0	31:8	Reserved (not overwritten)
	7:0	<b>[DevSNB] Barrier.Byte</b> If <b>Barrier Count Enable</b> was <u>set</u> on the barrier-completing BarrierMsg, this byte has a value of 0. If <b>Barrier Count Enable</b> was <u>clear</u> on the barrier-completing BarrierMsg, the value written is obtained from the Interface Descriptor. Format: U8



## 2. Unified Return Buffer (URB)

The Unified Return Buffer (URB) is a general-purpose buffer used for sending data between different threads, and, in some cases, between threads and fixed-function units (or vice-versa). A thread accesses the URB by sending messages.

### 2.1 URB Size

**[DevSNB]:** The GT1 product's URB provides 32KB of storage, arranged as 1024 256-bit rows. The GT2 product's URB provides 64KB of storage, arranged as 2048 256-bit rows. A row corresponds in size to an EU GRF register. Read/write access to the URB is generally supported on a row-granular basis.

A URB *entry* is a logical entity within the URB, referenced by an entry *handle* and comprised of some number of consecutive rows.

### 2.2 URB Access

The URB can be written by the following agents:

- Command Stream (CS) can write constant data into Constant URB Entries (CURBEs) as a result of processing CONSTANT\_BUFFER commands.
- The Video Front End (VFE) fixed-function unit of the Media pipeline can write thread payload data in to its URB entries.
- The Vertex Fetch (VF) fixed-function unit of the 3D pipeline can write vertex data into its URB entries
- Threads can write data into URB entries via URB\_WRITE messages sent to the URB shared function.

The URB can be read by the following agents:

- The Thread Dispatcher (TD) is the main source of URB reads. As a part of spawning a thread, pipeline fixed-functions provide the TD with a number of URB handles, read offsets, and lengths. The TD reads the specified data from the URB and provide that data in the thread payload pre-loaded into GRF registers.
- The Geometry Shader (GS) and Clipper (CLIP) fixed-function units of the 3D pipeline can read selected parts of URB entries to extract vertex data required by the pipeline.
- The Windower (WM) FF unit reads back depth coefficients from URB entries written by the Strip/Fan unit.

**[DevSNB]** Note that neither the CPU nor EU threads can read the URB directly.



## 2.3 State

The URB function is stateless, with all information required to perform a function being passed in the write message.

See URB Allocation (*Graphics Processing Engine*) for a discussion of how the URB is divided amongst the various fixed functions.

## 2.4 Messages [DevSNB]

There is only one type of message supported by the URB shared function: URB\_WRITE. It is primarily used by a thread to write data in to an entry in the URB, as referenced by the passed handle. FF units of the 3D pipeline snoop these messages, and a side effect of the message may be some information being passed to the FF unit which spawned the thread.

This section documents the global aspects of the URB write messages. The actual data contained in the message differs for each fixed function – refer to *3D Pipeline* and the fixed-function chapters or details on 3D URB data formats, *Media* for media-specific URB data formats, and *Graphics Processing Engine* for details on Constant URB Entries (CURBEs).

### Programming Notes:

- The **End of Thread** bit in the message descriptor may be set on URB messages only in threads dispatched by the vertex shader (VS), geometry shader (GS), clipper, and strips and fans (SF) units.
- The Kernel must do a write commit on the last write to DAP prior to a URB\_WRITE with **End of Thread**. [DevSNB]



## 2.4.1 Execution Mask

The Execution Mask specified in the 'send' instruction determines which DWords within each message register are written to the URB.

## 2.4.2 Message Descriptor

Bit	Description
19	<b>Header Present</b> This bit must be set to <b>one</b> for all URB messages.
18:16	Ignored
15	<b>Complete</b> If clear, this signals that the URB entry(s) referenced by the handle(s) are <u>not yet</u> completely specified. This setting is used to perform partial writes to URB entries, as would be required when writing an entry larger than the maximum single message payload can accommodate. Only the final write would be marked "complete". Partial writes may be unordered. If set, this signals that there will be no further writes (past this one) to the specific URB entry(s) by the thread. A snooping FF unit uses this to identify when the corresponding URB entry(s) are completely specified, at which point the FF unit can initiate further operations the entry(s) (either a readback, passing the handle(s) down the pipeline, or immediate deallocation if the entry is "unused"). This bit is strictly control information passed to snooping FF units. The URB shared function itself does not use this bit for any purpose. Programming Notes: <ul style="list-style-type: none"><li>• The following message descriptor fields are only valid when <b>Complete</b> is set: <b>Used</b></li><li>• The following message header fields are only valid when <b>Complete</b> is set: <b>Handle 0 PrimType, Handle 0 PrimStart, Handle 0 PrimEnd.</b></li></ul>



Bit	Description
14	<p><b>Used</b></p> <p>If set, this signals that the URB entry(s) referenced by the handle(s) are valid outputs of the thread. In all likelihood this means that that entry(s) contains complete &amp; valid data to be subject to further processing by the pipeline.</p> <p>If clear, this signals that the URB entry(s) referenced by the handle(s) are not valid outputs of the thread. Use of this setting will result in the handle(s) being immediately dereferenced by the owning FF unit. This setting is to be used by GS or CLIP threads to dereference handles it obtained (either in the initial thread payload or subsequent allocation writebacks) but subsequently determined were not required (e.g., the object was completely clipped out).</p> <p><b>Programming Notes:</b></p> <ul style="list-style-type: none"> <li>• <u>Only GS and CLIP threads are permitted to utilize Used==0.</u> All other threads are required (by design) to generate valid outputs in all cases.</li> <li>• This bit is strictly control information passed to snooping FF units. The URB shared function itself does not use this bit for any purpose.</li> <li>• This bit is only valid when <b>Complete</b> is set, i.e., it is ignored on partial writes.</li> </ul>
13	<p><b>Allocate</b></p> <p>If set, this requests that an additional destination URB entry be allocated to the thread by the spawning FF unit. The FF unit will return the handle to this URB entry via a message writeback operation in response to this message (see writeback format below). Therefore, threads <u>must</u> specify a writeback register in 'send' instructions issuing messages with this bit set.</p> <p>If clear, an additional allocation is not requested.</p> <p><b>Programming Notes:</b></p> <ul style="list-style-type: none"> <li>• This bit is strictly control information passed to snooping FF units. The URB shared function itself does not use this bit for any purpose.</li> <li>• This bit is valid on all URB_WRITE messages, e.g., it could be used to allocate a new handle on a partial write (<b>Complete</b> not set).</li> <li>• Only one Allocate request (per thread) can be outstanding. Upon requesting an allocation, the thread must wait for the handle to be returned (written back) before another allocation can be requested.</li> </ul>
12	<p><b>Fast Composite Restriction Check Pass</b></p> <p>If set on the end of thread message, this field indicates that the setup kernel portion of the fast composite restriction check has passed. This field is ignored for threads dispatched by units other than Strips and Fans. This field is also ignored unless at least one contiguous dispatch mode is enabled <i>and</i> at least one normal dispatch mode is enabled in WM_STATE.</p>
11:10	<p><b>Swizzle Control.</b> This field is used to specify which "swizzle" operation is to be performed on the write data. It indirectly specifies whether one or two handles are valid.</p> <p>00: URB_NOSWIZZLE The message data is to be written directly to a single URB entry (Handle 0).</p> <p>01: URB_INTERLEAVED The message contains data to be written to two URB entries. The message data provided is interleaved such that the upper DWords (7:4) of each 256-bit unit contain data to be written to Handle 1, and the lower DWords (3:0) contain data to be written to Handle 0. The URB shared function will de-interleave this data and write the two separate data streams to the two entries using the single Offset value (see Offset below for more details).</p> <p>10: URB_TRANSPOSE This message contains data that is to be "transposed" before being written to the URB. The transpose applied is tailored to the passing of data between the SF and WM stages – it is not a generic transpose. (See description below). Therefore, the assumption is that this mode will</p>



Bit	Description
	<p>only be used by Setup threads, where the setup-result data is swizzled before being written to the URB in order to provide a more optimal format for use in a subsequent PS thread. (See Strip/Fan, Windower chapters).</p> <p>If <b>Transposed URB Read Enable</b> (WM_STATE) is set, the Setup thread must use URB_NOSWIZZLE to write the coefficient data (it will be transposed whenever the URB is read). URB_TRANSPOSE must only be used when <b>Transposed URB Read Enable</b> is clear.</p> <p>See Programming Restrictions in the URB_TRANSPOSE subsection below.</p> <p>11: Reserved</p>
9:4	<p><b>Offset.</b> This field specifies a destination offset (in 256-bit units) from the start of the URB entry(s), as referenced by <b>URB Return Handle <i>n</i></b>, at which the data (if any) will be written.</p> <p>When URB_INTERLEAVED is used, this field provides a 256-bit granular offset applied to both URB entry destinations.</p> <p>When URB_TRANSPOSE is used, this field provides a 256-bit granular offset applied to the URB entry destination. The least significant bit of <b>Offset</b> must be zero.</p>
3:0	<p><b>URB Opcode</b></p> <p>0: URB_WRITE</p> <p>1: FF_SYNC</p> <p>all other codes are Reserved</p>

The following table lists the valid and invalid combinations of the Complete, Used, Allocate and EOT bits:

Complete	Used	Allocate	EOT	Valid?	Usage
0	d/c	0	0	Valid:	Normal partial-write or non-write of URB.
0	d/c	0	1	Valid only if any and all preceding URB entries have been marked as "complete" and there is no outstanding Allocate request.	Thread terminate w/ non-write of URB
0	d/c	1	0	Valid only if any and all preceding URB entries have been marked as "complete" and there is no outstanding Allocate request.	Non-write of URB with request for an additional handle.
0/1	d/c	1	1	Invalid. Thread must never terminate with an outstanding writeback request.	n/a
1	0	0/1	0	Valid	Dereference of URB entry without/with new allocation request.
1	0	0	1	Valid	Dereference of URB entry and thread termination.
1	1	0/1	0	Valid	Completion of URB entry output without/with new allocation request.



Complete	Used	Allocate	EOT	Valid?	Usage
1	1	0	1	Valid	Completion of URB entry output and thread termination.

## 2.4.3 URB\_WRITE

### 2.4.3.1 URB\_WRITE Message Header

DWord	Bit	Description
M0.7	31:0	<b>Reserved</b>
M0.6	31:0	<b>Reserved</b>
M0.5	31:8	Ignored
	7:0	<b>FFTID.</b> The Fixed Function Thread ID is assigned by the fixed function unit and is a unique identifier for the thread. It is used to free up resources used by the thread upon thread completion.
M0.4	31:0	Ignored
M0.3	31:0	Ignored
M0.2	31:27	Ignored
	26:16	<b>SONumPrimsWritten Increment Value.</b> <b>[DevSNB+]</b> This field is ignored if the <b>SVBI Post Increment</b> field is set in the GS_3DSTATE in the 3D pipeline chapter.
	15:8	Ignored
	7	Ignored
	6:2	<b>Handle 0 PrimType.</b> This field associates a primitive type with the vertex written at Handle 0. <b>NOTE:</b> This field is only defined when the GS or Clipper FF unit is the target FF unit. Otherwise it is Reserved:MBZ.
1	1	<b>Handle 0 PrimStart.</b> This field is used to indicate that the vertex written at Handle 0 is the <u>first</u> vertex of a primitive. <b>NOTE:</b> This field is only defined when the GS or Clipper FF unit is the target FF unit. Otherwise it is Reserved:MBZ.
	0	<b>Handle 0 PrimEnd.</b> This field is used to indicate that the vertex written at Handle 0 is the <u>last</u> vertex of a primitive. <b>NOTE:</b> This field is only defined when the GS or Clipper FF unit is the target FF unit. Otherwise it is Reserved:MBZ.
M0.1	31:16	<b>Handle ID 1.</b> This ID is assigned by the fixed function unit and links the work in channel 1 to a specific entry within the fixed function unit. This field is ignored unless <b>Swizzle Control</b> indicates Interleave mode.
	15:0	<b>URB Return Handle 1.</b> This is the URB handle where channel 1's results are to be placed. This field is ignored unless <b>Swizzle Control</b> indicates interleave mode.



DWord	Bit	Description
M0.0	31:16	<b>Handle ID 0.</b> This ID is assigned by the fixed function unit and links the work in channel 0 to a specific entry within the fixed function unit.
	15:0	<b>URB Return Handle 0.</b> This is the URB handle where channel 0's results are to be placed.

### 2.4.3.2 URB\_WRITE Message Payload

For the URB message, the message payload will be written to the URB entries indicated by the URB return handles in the message header.

While GS and CLIP threads will write one vertex at a time to the URB, the VS will write two interleaved vertices. The description of the URB write messages will refer to the per-vertex DWords described in the Vertex URB Entry Formats section of the *3D Overview* chapter.

Payload	Usage
URB_NOSWIZZLE	The message payload contains data to be written to a single URB entry (e.g., one Vertex URB entry). The <b>Swizzle Control</b> field of the message descriptor must be set to 'NoSwizzle'.
URB_INTERLEAVED	The message payload contains data to be written to two separate URB entries. The payload data is provided in a high/low interleaved fashion. The <b>Swizzle Control</b> field of the message descriptor must be set to 'Interleave'.
URB_TRANSPOSE	The message payload contains data that is to be transposed before being written to the URB. See the <i>Strip &amp; Fan (SF) Unit</i> chapter for details on the source and destination data layouts and intended usage model.





### 2.4.3.2.1 URB\_NOSWIZZLE

URB\_NOSWIZZLE is used to simply write data into consecutive URB locations (no data swizzling or transposition applied).

#### Programming Notes:

- The URB function will ignore the Channel Enables associated with this message and write all channels into the URB.
- When **Transposed URB Read Enable** (WM\_STATE) is set, the Setup thread must use URB\_NOSWIZZLE to write the coefficient data into the URB (it will be transposed whenever the URB is read).

When URB\_NOSWIZZLE is used to write vertex data, the following table shows an example layout of a URB\_NOSWIZZLE payload containing one (non-interleaved) vertex containing  $n$  pairs of 4-DWord vertex elements (where for the example,  $n$  is  $>2$ ).

DWord	Bit	Description
M1.7	31:0	Vertex Data [7]
M1.6	31:0	Vertex Data [6]
M1.5	31:0	Vertex Data [5]
M1.4	31:0	Vertex Data [4]
M1.3	31:0	Vertex Data [3]
M1.2	31:0	Vertex Data [2]
M1.1	31:0	Vertex Data [1]
M1.0	31:0	Vertex Data [0]
M2.7	31:0	Vertex Data [15]
M2.6	31:0	Vertex Data [14]
M2.5	31:0	Vertex Data [13]
M2.4	31:0	Vertex Data [12]
M2.3	31:0	Vertex Data [11]
M2.2	31:0	Vertex Data [10]
M2.1	31:0	Vertex Data [9]
M2.0	31:0	Vertex Data [8]
...		...
Mn.7	31:0	Vertex Data $[8(n-1)+7]$
Mn.6	31:0	Vertex Data $[8(n-1)+6]$
Mn.5	31:0	Vertex Data $[8(n-1)+5]$
Mn.4	31:0	Vertex Data $[8(n-1)+4]$
Mn.3	31:0	Vertex Data $[8(n-1)+3]$
Mn.2	31:0	Vertex Data $[8(n-1)+2]$



DWord	Bit	Description
Mn.1	31:0	Vertex Data [8(n-1)+1]
Mn.0	31:0	Vertex Data [8(n-1)+0]

### 2.4.3.2.2 URB\_INTERLEAVED

The following table shows an example layout of a URB\_INTERLEAVED payload containing two interleaved vertices, each containing  $n$  4-DWord vertex elements ( $n > 1$ ).

#### Programming Restrictions:

- At least 256 bits per vertex (512 bits total, M1 & M2) must be written. Writing only 128 bits per vertex (256 bits total, M1 only) results in UNDEFINED operation.
- The URB function will use (not ignore) the Channel Enables associated with this message.
- **[DevSNB]** Interleave writes must be in multiples of 256 per vertex.

DWord	Bit	Description
M1.7	31:0	Vertex 1 Data [3]
M1.6	31:0	Vertex 1 Data [2]
M1.5	31:0	Vertex 1 Data [1]
M1.4	31:0	Vertex 1 Data [0]
M1.3	31:0	Vertex 0 Data [3]
M1.2	31:0	Vertex 0 Data [2]
M1.1	31:0	Vertex 0 Data [1]
M1.0	31:0	Vertex 0 Data [0]
M2.7	31:0	Vertex 1 Data [7]
M2.6	31:0	Vertex 1 Data [6]
M2.5	31:0	Vertex 1 Data [5]
M2.4	31:0	Vertex 1 Data [4]
M2.3	31:0	Vertex 0 Data [7]
M2.2	31:0	Vertex 0 Data [6]
M2.1	31:0	Vertex 0 Data [5]
M2.0	31:0	Vertex 0 Data [4]
...		...
Mn.7	31:0	Vertex 1 Data [4(n-1)+3]
Mn.6	31:0	Vertex 1 Data [4(n-1)+2]
Mn.5	31:0	Vertex 1 Data [4(n-1)+1]
Mn.4	31:0	Vertex 1 Data [4(n-1)+0]



DWord	Bit	Description
Mn.3	31:0	Vertex 0 Data [4(n-1)+3]
Mn.2	31:0	Vertex 0 Data [4(n-1)+2]
Mn.1	31:0	Vertex 0 Data [4(n-1)+1]
Mn.0	31:0	Vertex 0 Data [4(n-1)+0]

### 2.4.3.2.3 URB\_TRANSPOSE

The following table shows an example layout of a URB\_TRANSPOSE payload and how the data is transposed and stored in the destination URB entry. Note that Source Row 0, Source Row 1, and implied row of all-zero, and Source Row 3 is transposed and stored in successive 4-DW locations in the destination. This is then repeated for the next 3 rows of the source payload. For the intended usage model in the Setup thread, Source Row 0 would contain “Cx” coefficients for the first 8 attributes, Source Row 1 would contain “Cy” coefficients for the first 8 attributes, and Source Row 2 would contain “C0” coefficients for the first 8 attributes, then repeating for the next 8 attributes. Insertion of the implied all-zero row is required to align the Cx,Cy and C0 attributes into half-rows within the URB. This permits the used of the “LINE” instruction to initiate attribute interpolation in the subsequent PS thread.

#### Programming Notes:

- The message payload must contain a multiple of 3 Source Rows of data (excluding the message header).
- The URB function will ignore the Channel Enables associated with this message and write all channels into the URB.
- When **Transposed URB Read Enable** (WM\_STATE) is set, the Setup thread must use URB\_NOSWIZZLE to write the coefficient data into the URB (it will be transposed whenever the URB is read). URB\_TRANSPOSE should only be used when **Transposed URB Read Enable** is clear.

**Table 2-1. URB\_TRANSPOSE Payload**

DWord	Bit	Description
M1.0-7	31:0	Source Row 0 (e.g., Cx coeffs for the 1 <sup>st</sup> set of 8 attributes)
M2.0-7	31:0	Source Row 1 (e.g., Cy coeffs for the 1 <sup>st</sup> set of 8 attributes)
M3.0-7	31:0	Source Row 2 (e.g., C0 coeffs for the 1 <sup>st</sup> set of 8 attributes)
M4.0-7	31:0	Source Row 3 (e.g., Cx coeffs for the 2 <sup>nd</sup> set of 8 attributes)
M5.0-7	31:0	Source Row 4 (e.g., Cy coeffs for the 2 <sup>nd</sup> set of 8 attributes)
M6.0-7	31:0	Source Row 5 (e.g., C0 coeffs for the 2 <sup>nd</sup> set of 8 attributes)
...	31:0	...



**Table 2-2. URB\_TRANSPOSE URB Destination Layout**

	URB DW							
URB Row	7	6	5	4	3	2	1	0
n+0	M3.1	0	M2.1	M1.1	M3.0	0	M2.0	M1.0
n+1	M3.3	0	M2.3	M1.3	M3.2	0	M2.2	M1.2
n+2	M3.5	0	M2.5	M1.5	M3.4	0	M2.4	M1.4
n+3	M3.7	0	M2.7	M1.7	M3.6	0	M2.6	M1.6
n+4	M6.1	0	M5.1	M4.1	M6.0	0	M5.0	M4.0
n+5	M6.3	0	M5.3	M4.3	M6.2	0	M5.2	M4.2
n+6	M6.5	0	M5.5	M4.5	M6.4	0	M5.4	M4.4
n+7	M6.7	0	M5.7	M4.7	M6.6	0	M5.6	M4.6
	...	...	...	...	...	...	...	...

### 2.4.3.3 Writeback Message for URB Entry Allocate

A writeback only occurs if the **Allocate** bit is set in the message descriptor. A single register is returned containing the URB Return Handle and Handle ID for the allocated handle in the low DWord is returned. All high DWords contain zero.

DWord	Bit	Description
W0.7:1		Reserved : MBZ
W0.0	31:16	<b>Handle ID.</b> This ID is assigned by the fixed function unit and links the thread to a specific entry within the fixed function unit.
	15:0	<b>URB Return Handle.</b> This is the URB handle where the thread's results are to be placed.

### 2.4.4 FF\_SYNC Messages

The FF\_SYNC message is used to pass critical information between GS/Clip threads and the GS/Clip FF units, as well as providing Gs/Clip thread synchronization (ordering). GS threads report various counts resulting from running the GS and/or SO functions, prior to performing any output (to SOB buffers or to URB handles). Clip threads report only # of handles required. A message response (writeback) length of 1 GRF will be indicated on the 'send' instruction if the thread requires response data and/or synchronization. Refer to the GS/Clip stage chapter for details.



### 2.4.4.1 FF\_SYNC Message Header

Dword	Bit	Description
M0.7	31:0	<b>Reserved</b>
M0.6	31:0	<b>Reserved</b>
M0.5	31:8	Ignored
	7:0	<b>FFTID.</b> This ID is assigned by the fixed function unit and is a unique identifier for the thread. It is used to free up resources used by the thread upon thread completion.
M0.4	31:0	Ignored
M0.3	31:0	Ignored
M0.2	31:7	Ignored
M0.1	31:16	Ignored
	15:0	<b>(GS-only) NumGSPrimsGenerated.</b> The number of objects (e.g., triangles) generated by the GS function performed by the thread. If the GS function is not enabled, this field MBZ. Format: U16 Range: [0,1024]
M0.0	31:16	<b>(GS-only) NumSOVertsToWrite.</b> The number of (expanded-to-list) vertices generated by the SO function performed by the thread. This represents the number of vertices the thread will attempt to write to the SOB(s) in memory, once it obtains the SVBI(s) in the FF_SYNC writeback. Note that overflow may occur either (a) prior to the SVBI(s) are returned in the writeback or (b) in the process of this thread outputting to the SOB(s). In either case, the thread needs to check for overflow once it receives the writeback, based on the returned SVBI(s) and the number of vertices it must attempt to output. If the SO function is not enabled, this field MBZ. <b>[DevSNB+]</b> This field is ignored if the <b>SVBI Post Increment</b> field is set in the GS_3DSTATE in the 3D pipeline chapter. Format: U16 Range: [0,3066] (1024-vertex trisrip = 1022 triangles = 3066 trillist vertices)
	15:0	<b>(GS-only) NumSOPrimsNeeded.</b> The number of objects (e.g., triangles within a trillist) generated by the SO function performed by the thread (exclusive of any SOB overflow). If the SO function is not enabled, this field MBZ. <b>[DevSNB+]</b> This field is ignored if the <b>SVBI Post Increment</b> field is set in the GS_3DSTATE in the 3D pipeline chapter. Format: U16 Range: [0,1024]

### 2.4.4.2 FF\_SYNC Writeback Message

(Both GS & Clip): DWord W0.0 of the writeback data contains initial handle information. If **Handle Valid** is clear, the FF unit did not have a handle available to be allocated as the initial handle – the thread will need to use **the URB\_WRITE message to obtain the initial handle. Otherwise the Handle ID and URB Return Handle** fields are valid and can be used to write the first VUE.

(GS-only) The writeback data contains the SVBI values to be used as starting write indices by the GS thread. It is the responsibility of the GS thread to perform SOB overflow processing. If the GS thread is



not performing StreamOutput and was simply using the writeback to provide GS vertex output synchronization, the return data is to be ignored.

(Clip-only) Dwords W0.1-7 of the writeback data are ignored.

DWord	Bit	Description
W0.7: W0.5	31:0	Reserved
W0.4	31:0	<b>(GS-only) Streamed Vertex Buffer Index 3</b> This field represents the value of SVBI[3] that is to be used as the starting index for the GS thread. If the thread is not performing StreamOutput, this field is ignored. Format = U32 Range = $[0, 2^{27}-1]$
W0.3	31:0	<b>(GS-only) Streamed Vertex Buffer Index 2</b>
W0.2	31:0	<b>(GS-only) Streamed Vertex Buffer Index 1</b>
W0.1	31:0	<b>(GS-only) Streamed Vertex Buffer Index 0</b>
W0.0	31:24	Reserved
	23:16	<b>Handle ID.</b> This ID is assigned by the FF unit and links the thread to a specific entry within the FF unit. Format: <u>Reserved for Implementation Use</u>
	15:12	Reserved
	11:0	<b>URB Return Handle.</b> This is the initial destination URB handle passed to the thread. If the thread does output URB entries, this identifies the first destination URB entry. <b>[DevSNB+]</b> Format: U12 256 bit URB Handle Address



## 3. Execution Unit ISA

### 3.1 Introduction

#### 3.1.1 Objective and Scope

The core architecture consists of an array of multi-threaded processors, also referred to as Execution Units (EU). This Instruction Set Architecture (ISA) document specifies the instructions executable on the EUs of the architecture. It defines the data types in the GEN architecture. It includes the binary format (machine code) and ASCII format (native syntax) of each instruction. It also provides example usages of instructions and modes of instructions, and certain data formats. The programming guideline in appendix provides information to help developers to understand the usage of ISA. However, it is not intended to be a comprehensive tutorial.

#### 3.1.2 Terms and Acronyms

AIP	Application IP. This is part of the control registers for exception handling for a thread. Upon an exception, hardware moves the current IP into this register and then jumps to SIP.
ARF	Architecture Register File. It is a collection of architecturally visible registers for a thread such as address registers, accumulator, flags, notification registers, IP, null, etc. ARF should not be mistaken as just the address registers.
B	Byte. As a numerical data type of 8 bits, B represents a signed byte integer. It is used to specify the type of an operand in an instruction.
BNF	Backus Naur Form, a formal notation to describe the syntax of a given language. The meta symbols of BNF include “::=”, “ ”, and “< >”, where “::=” means “is defined as”; “ ” means “or”; and angle brackets “<” and “>” are used to surround category names.
CR	Control Register. These read-write registers are used for thread mode control and exception handling for a thread.
D	Double word (DWord). As a fundamental data type, D or DW represents 4 bytes. It may be used to specify the type of an operand in an instruction.
EOT	End Of Thread. This is a message sideband signal on the Output message bus signifying that the message requester thread is terminated. A thread must have at least one SEND instruction with the EOT bit in the message descriptor field set in order to properly terminate.
EU	Execution Unit. An EU is a multi-threaded processor within the multi-processor system. Each EU is a fully-capable processor containing instruction fetch and decode, register files, source operand swizzle and SIMD ALU, etc. An EU is also referred to as a Core.



EUID	Execution Unit Identifier. The 4-bit field within a thread state register (SR0) that identifies the row and column location of the EU where a thread is located. A thread can be uniquely identified by the EUID and TID.
ExecSize	Execution Size.
Execution Size	Execution Size indicates the number of data elements processed by a SIMD instruction. It is one instruction field and can be changed at a per instruction level.
FLT_MAX	The magnitude of the maximum represent-able single-precision floating number according to IEEE-754 standard. FLT_MAX has an exponent of 0xFE and a mantissa of all one's.
fmax	Same as FLT_MAX.
Core	Alternative name for an EU in the multi-processor system.
GRF	General Register File. This is the most commonly used read-write register space organized as an array of 256-bit registers for a thread.
ISA	Instruction Set Architecture. The ISA describes the instructions supported by an EU. A sequence of instructions forms a thread executed on an EU.
JIT	Just-In-Time compiler
LSB	Least Significant Bit
Message	Messages are data packages transmitted from a thread to another thread, to another shared function or to another fixed function. Message passing is the primary communication mechanism of the architecture.
MRF	Message Register File. This is the write-only register space, organized as an array of 256-bit registers, for a thread to communicate with shared functions or other threads.
MSB	Most Significant Bit
DQ	Double Quad word (DQword). As a fundamental data type, DQ represents 16 bytes.
POR	Plan Of Record
QW	Quad Word (QWord). As a fundamental data type, QW represents 8 bytes.
QQ	Quad Quad word (QQword). As a fundamental data type, QQ represents 32 bytes.
Sub-Register	Subfield of a SIMD register. A SIMD register is an aligned fixed size register for a register file or a register type. For example, a GRF register, <i>r2</i> , is a 256-bit wide, 256-bit aligned register. A sub-register, <i>r2.3:d</i> , is the fourth dword of GRF register <i>r2</i> .
SIMD	Single Instruction Multiple Data. The term SIMD can be used to describe the kind of parallel processing architecture that exploits data parallelism at the instruction level. It can also be used to describe the instructions in such an architecture.



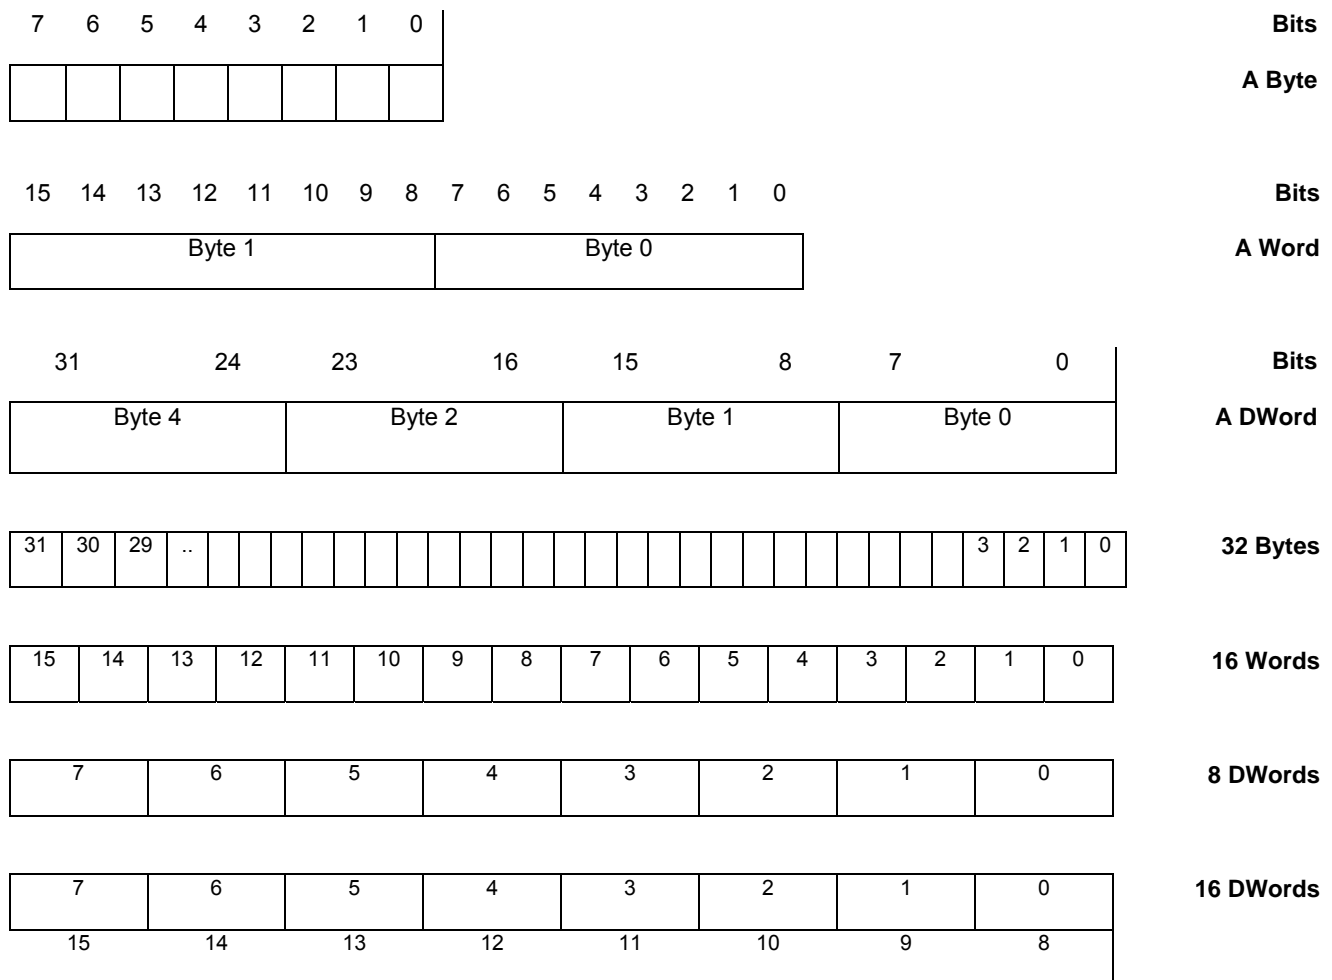


SIP	System IP. There is one global System IP register for all the threads. From a thread's point of view, this is a virtual read-only register. Upon an exception, hardware performs certain book-keeping functions and then jumps to SIP.
SR	State Register. The read-only registers containing the state information of the current thread, including the EUID/TID, Dispatcher Mask, and System IP.
Thread	A thread is an instance of a kernel program executed on an EU. The life cycle for a thread starts from the executing the first instruction after being dispatched from Thread Dispatcher to an EU to the execution of the last instruction – a send instruction with EOT that signals the thread termination. Threads in the system may be independent from each other or communicate with each other through the Message Gateway share function.
TID	Thread Identifier. The 2-bit field within a thread state register (SR0) that identifies which out of the four possible thread slots on the EU is executing that thread. A thread can be uniquely identified by the EUID and TID.
TS	Thread Spawner. TS is the second and the last fixed function stage of the media pipeline.
V	Immediate integer vector. As a numerical data type of 32 bits, an immediate integer vector of type V contains 8 signed integer elements with 4 bits each. The 4-bit integer element is in 2's complement form. It may be used to specify the type of an immediate operand in an instruction.
VF	Immediate floating point vector. As a numerical data type of 32 bits, an immediate floating point vector of type VF contains 4 floating point elements with 8-bit each. The 8-bit floating point element contains a sign field, a 3-bit exponent field and a 4-bit mantissa field. It may be used to specify the type of an immediate operand in an instruction.
W	Word. As a numerical data type of 16 bits, W represents a signed word integer. It is used to specify the type of an operand in an instruction.
URB	Unified Return Buffer. The on-chip memory managed/shared by Fixed Functions. Threads use the URB to return data that will be consumed either by a Fixed Function or other threads.
UB	Unsigned Byte integer. A numerical data type of 8 bits. It may be used to specify the type of an operand in an instruction.
UD	Unsigned Double Word integer. A numerical data type of 32 bits. It may be used to specify the type of an operand in an instruction.
UW	Unsigned Word integer. A numerical data type of 16 bits. It may be used to specify the type of an operand in an instruction.
VFE	Video Front End. VFE is the first fixed function stage of the media pipeline.



### 3.1.3 Formats and Conventions

In order to conveniently (and without ambiguity) describe the register files with 256-bit wide registers that may contain various data types with different data element widths, it is important to use a consistent table format to represent the registers. Throughout this document, we will adopt the following table formats and conventions. When a register or a number is presented by a row, increasing order is always **from right to left** and then **top down** pictorially. In other words, for a bit field, the LSB to MSB is from right to left; for a byte sequence, the least significant byte to the most significant byte is also from right to left. This is consistent with the 'Little Endian' convention used by IA-32 machines. The following tables depict the layout formats for different data units.



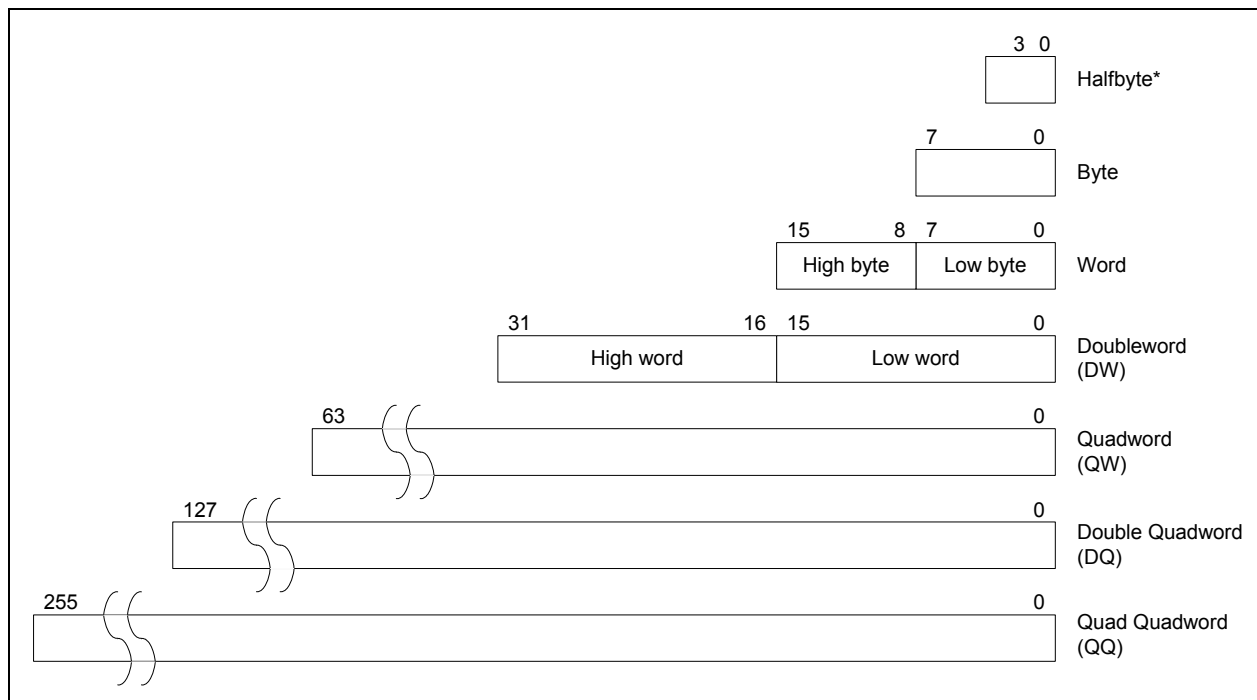
With this convention, we note that the execution channels are logically viewed as from right to left too, which is a little bit unconventional. However, as shown in the *GEN Execution Environment* Chapter, it matches with the bit order of the flag registers. This also impacts the view of a GRF register region, now the region origin is located at the upper-right corner and a region row is viewed from right to left.

## 4. EU Data Types

### 4.1 Fundamental Data Types

The fundamental data types in the architecture are halfbyte, byte, word, doubleword (DW), quadword (QW), double quadword (DQ) and quad quadword (QQ). They are defined based on the number of bits of the data type, ranging from 4 bits to 256 bits. As shown in **Error! Reference source not found.**, a halfbyte contains 4 bits, a byte contains 8 bits, a word contains two bytes, and a doubleword (dword) contains two words, and so on. Halfbyte is a special data type such that it cannot be accessed directly as standalone data element. It is only allowed as a subfield of the numerical data type of “packed signed halfbyte integer vector” described in the next section.

Figure 4-1 Fundamental data types



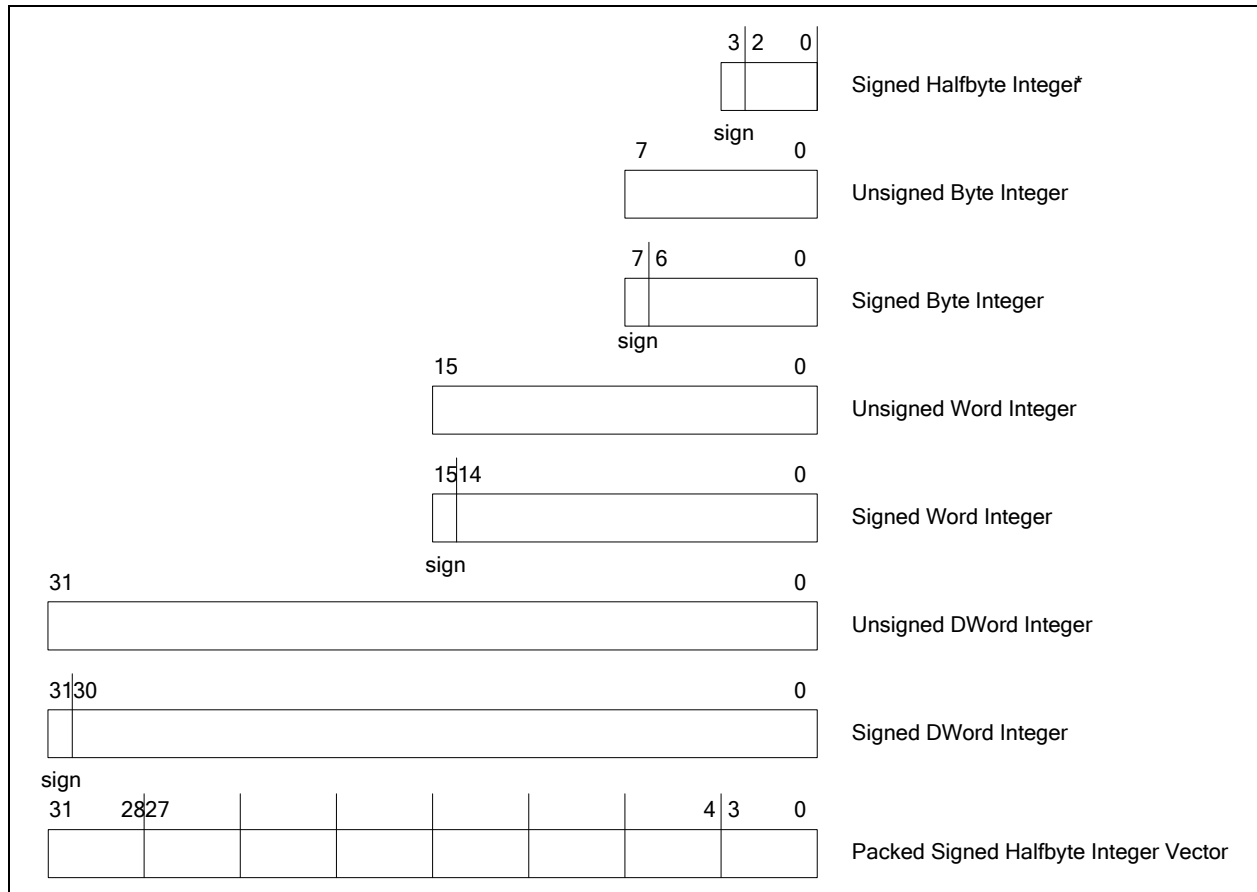
With the exception of halfbyte, the access of a data element to/from a register or to/from memory must be aligned on the natural boundaries of the data type. The natural boundary for a word has an even-numbered address in unit of byte. The natural boundary for a doubleword has an address divisible by 4 bytes. Similarly, the natural boundary for a quadword, double quadword and quad quadword has an address divisible by 8, 16, and 32 bytes, respectively. Quadword, double quadword and quad quadword do not have corresponding numerical data type. Instead, they are used to describe a group (a vector) of numerical data elements of smaller size align to larger natural boundaries.



## 4.2 Numerical Data Types


The numerical data types defined in the architecture include signed and unsigned integers and floating-point numbers (floats) of various numbers of bits. The following illustration details the notation, size and numerical range of each data type. The largest numerical data type has 32 bits.

**Figure 4-2 Integer numerical data types**





**Figure 4-3. Floating point numerical data types**



**Table 4-1. Formats and ranges of numerical data types**

Notation	Numerical Data Types	Fundamental Data Type	Range
UB	Unsigned Byte Integer	Byte	[0, 255]
B	Signed Byte Integer	Byte	[-128, 127]
UW	Unsigned Word Integer	Word	[0, 65535]
W	Signed Word Integer	Word	[-32768, 32767]
UD	Unsigned Doubleword Integer	Doubleword	[0, $2^{32} - 1$ ]
D	Signed Doubleword Integer	Doubleword	[- $2^{31}$ , $2^{31} - 1$ ]
F	Single Precision Float	Doubleword	[-( $2^{-23}$ ) <sup>127</sup> ...- $2^{-149}$ , 0.0, $2^{-149}$ ... ( $2^{-23}$ ) <sup>127</sup> ]
n/a	Signed Halfbyte Integer	Halfbyte	[-8, 7]
V	Packed Signed Halfbyte Integer Vector	Doubleword	[-8, 7]
n/a	Restricted 8-bit Float	Byte	[-31...-0.125, 0, 0.125... 31]
VF	Packed Restricted Float Vector	Doubleword	[-31...-0.125, 0, 0.125... 31]



### 4.2.1 Unsigned Integers

Unsigned integers are unsigned binary numbers contained in a byte, a word or a doubleword. The range for an unsigned byte integer is from 0 to 255. The range for an unsigned word integer is from 0 to 65535. The range for an unsigned doubleword integer is from 0 to  $2^{32} - 1$ .

The short hand notation for an unsigned byte integer, an unsigned word integer, and an unsigned doubleword integer is **UB**, **UW**, **UD**, respectively.

### 4.2.2 Signed Integers

Signed integers are signed binary number in 2's compliment form contained in a halfbyte, a byte, a word or a doubleword. A signed halfbyte integer has a numerical range from -8 to 7 with the sign bit at bit 3. A signed byte integer has a range from -128 to 127 with the sign at bit 7. A signed word integer is has a range from -32768 to 32767 with the sign at bit 15. A signed doubleword integer has a range from  $-2^{31}$  to  $2^{31} - 1$  with the sign at bit 31.

The short hand notation for a signed byte integer, a signed word integer, and a signed doubleword integer is **B**, **W**, **D**, respectively.

### 4.2.3 Single Precision Floating-Point Numbers

The single precision floating point numbers is contained in a doubleword. Floating point format is as defined in IEEE Standard 754 for Binary Floating-Point Arithmetic. Maximal representable number is  $(2-2^{-23})^{127}$  and the minimal number is  $-(2-2^{-23})^{127}$ . The smallest fractional negative number  $-2^{-149}$  and the smallest fractional positive number is  $2^{-149}$ . Value 0.0 has no fractional parts.


The short hand format notation for a single precision floating-point number is **F**.

### 4.2.4 Packed Signed Half-Byte Integer Vector

A packed signed halfbyte integer vector consists of 8 signed halfbyte integers contained in a doubleword. Each signed halfbyte integer element has a range from -8 to 7 with the sign on bit 3. This numerical data type is only used by an immediate source operand of doubleword in an instruction. It cannot be used for the destination operand or a non-immediate source operand. The hardware converts the 32-bit vector into 8-element signed word vector by sign extension. This is illustrated in Figure 4-4.

The short hand format notation for a packed signed half-byte vector is **V**.

**Figure 4-4. Converting a packed half-byte vector to a 128-bit signed integer vector**



### 4.2.5 Packed 8-bit Restricted Float Vector

A packed restricted float vector consists of 4 8-bit restricted floats contained in a doubleword. Each restricted float has the sign at bit 7, a 3-bit coded exponent in bits 4 to 6, a 4-bit fraction in bits 0 to 3, and an implied integer 1. The exponent is in excess-3 format – having a bias of 3. Restricted float provides zero, positive/negative normalized numbers with a small range (3-bit exponent) and small precision (4-bit fraction). This numerical data type is only used by an immediate source operand of doubleword in an instruction. It cannot be used for the destination operand, or a non-immediate source operand.

Figure 4-5 shows how to convert an 8-bit restricted float into a single precision float. Converting a 3-bit exponent with a bias of 3 to an 8-bit exponent with a bias of 127 is by adding 4, or equivalently copying bit 2 to bit 7 and putting the inverted bit 2 to bits 6:2. A special logic is also needed to take care of positive/negative zeros.



**Figure 4-5. Conversion from a Restricted 8-bit Float to a Single-Precision Float**

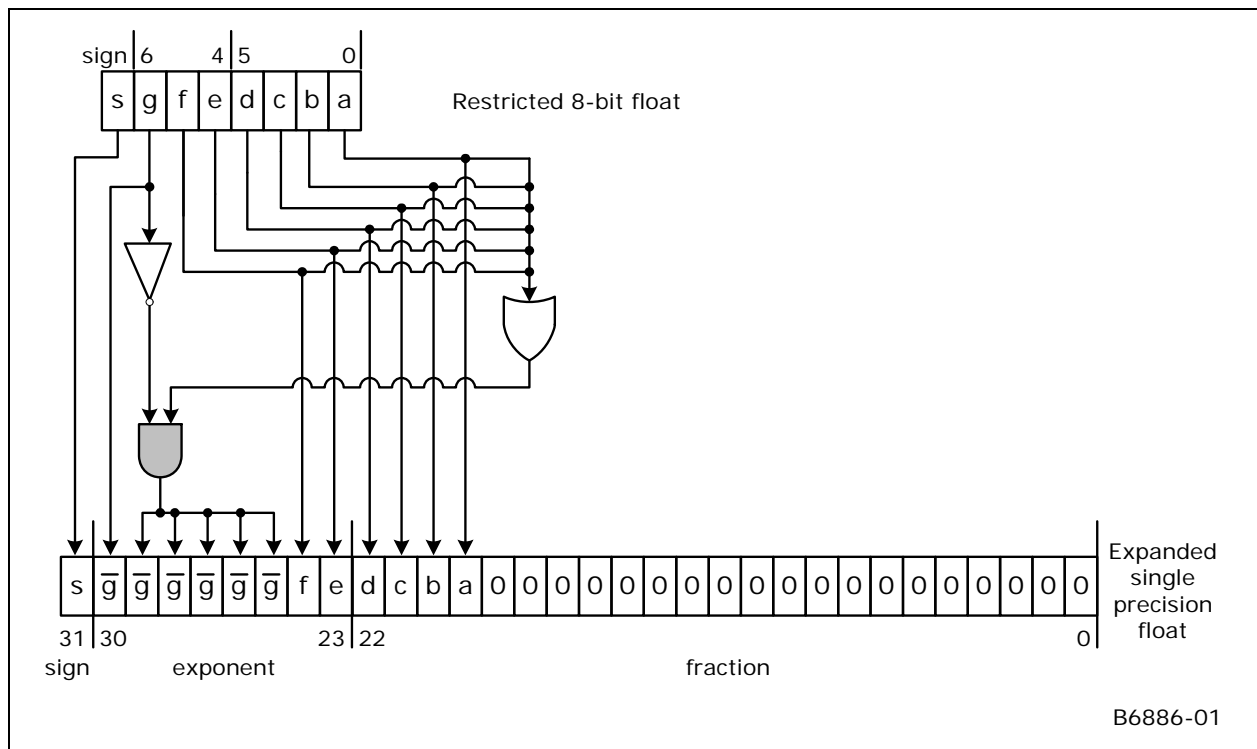


Table 4-2 shows all possible numbers of the restricted 8-bit float. Only normalized float numbers can be represented, including positive and negative zero, and positive and negative finite numbers. Normalized infinities, NaN and denormalized float numbers cannot be represented by this type. It should be noted that this 8-bit floating point format does not follow IEEE-754 convention in describing numbers with small magnitudes. Specifically, when the exponent field is zero and the fraction field is not zero, an implied one is still present instead of taking a denormalized form (without an implied one). This results in a simple implementation but with a smaller dynamic range – the magnitude of the smallest non-zero number is 0.125.






**Table 4-2. Example of restricted 8-bit float numbers**

Class	Restricted 8-bit Float				Extended 8-bit Exponent	Floating number in decimal
	Hex #	Sign [7]	Exponent [6:4]	Fraction [3:0]		
Positive Normalized Float	0x70-0x7F	0	<b>111</b>	0000 ... 1111	<b>1000 0011</b>	16 ... 31
	0x60-0x6F	0	<b>110</b>	0000 ... 1111	<b>1000 0010</b>	8 ... 15.5
	0x50-0x5F	0	<b>101</b>	0000 ... 1111	<b>1000 0001</b>	4 ... 7.75
	0x40-0x4F	0	<b>100</b>	0000 ... 1111	<b>1000 0000</b>	2 ... 3.875
	0x30-0x3F	0	<b>011</b>	0000 ... 1111	<b>0111 1111</b>	1 ... 1.9375
	0x20-0x2F	0	<b>010</b>	0000 ... 1111	<b>0111 1110</b>	0.5 ... 0.96875
	0x10-0x1F	0	<b>001</b>	0000 ... 1111	<b>0111 1101</b>	0.25 ... 0.484375
	0x01-0x0F	0	<b>000</b>	0001 ... 1111	<b>0111 1100</b>	0.125 ... 0.2421875
	0x00	0	<b>000</b>	<b>0000</b>	0000 0000	0 (+zero)
Negative Normalized Float	0xF0-0xFF	1	<b>111</b>	0000 ... 1111	<b>1000 0011</b>	-16 ... -31
	0xE0-0xEF	1	<b>110</b>	0000 ... 1111	<b>1000 0010</b>	-8 ... -15.5
	0xD0-0xDF	1	<b>101</b>	0000 ... 1111	<b>1000 0001</b>	-4 ... -7.75
	0xC0-0xCF	1	<b>100</b>	0000 ... 1111	<b>1000 0000</b>	-2 ... -3.875
	0xB0-0xBF	1	<b>011</b>	0000 ... 1111	<b>0111 1111</b>	-1 ... -1.9375
	0xA0-0xAF	1	<b>010</b>	0000 ... 1111	<b>0111 1110</b>	-0.5 ... -0.96875
	0x90-0x9F	1	<b>001</b>	0000 ... 1111	<b>0111 1101</b>	-0.25 ... -0.484375
	0x81-0x8F	1	<b>000</b>	0001 ... 1111	<b>0111 1100</b>	-0.125 ... -0.2421875
	0x80	1	<b>000</b>	<b>0000</b>	0000 0000	-0 (-zero)

Figure 4-6 shows the conversion of a packed exponent-only float to a 4-element vector of single precision floats.

The shorthand format notation for a packed signed half-byte vector is **VF**.

**Figure 4-6. Converting a Packed Restricted Float Vector to a 128-bit Float Vector**



### 4.3 Floating Point Modes

The architecture supports two floating point operation modes, namely IEEE floating point mode (IEEE mode) and alternative floating point mode (ALT mode). Both modes follow mostly the requirements in IEEE-754 but with different deviations. The deviations will be described in details in later sections. The primary difference between these modes is on the handling of Infs, NaNs and denorms. The IEEE floating point mode may be used to support newer versions of 3D graphics API Shaders and the alternative floating point mode may be used to support early Shader versions.

These two modes are supported by all units that perform floating point computations, including execution units, shared functions like Extended Math, the Sampler and the Render Cache color calculator, and fixed functions like VF, Clipper, SF and WIZ. Host software sets floating point mode through the fixed function state descriptors for 3D pipeline and the interface descriptor for media pipeline. Therefore different modes may be associated with different threads running concurrently. Floating point mode control for EU and shared functions are based on the floating point mode field (bit 0) of *cr0* register.

#### 4.3.1 IEEE Floating Point Mode

##### 4.3.1.1 Partial Listing of Honored IEEE-754 Rules

Here is a summary of expected 32-bit floating point behaviors in the architecture. Refer to IEEE-754 for topics not mentioned.

- $INF - INF = NaN$
- $0 * (+/-)INF = NaN$
- $1 / (+INF) = +0$  and  $1 / (-INF) = -0$ 
  - $(+/-)INF / (+/-)INF = NaN$  as  $A/B = A * (1/B)$
- $INV (+0) = RSQ (+0) = +INF$ ,  $INV (-0) = RSQ (-0) = -INF$ , and  $SQRT (-0) = -0$
- $RSQ (-finite) = SQRT (-finite) = NaN$
- $LOG (+0) = LOG (-0) = -INF$ ,  $LOG (-finite) = LOG (-INF) = NaN$



- NaN (any OP) any-value = NaN with one exception for min/max mentioned below. Resulting NaN may have different bit pattern than the source NaN.
- Normal comparison with conditional modifier of EQ, GT, GE, LT, LE, when either or both operands is NaN, returns FALSE. Normal comparison of NE, when either or both operands is NaN, returns TRUE.
  - Note: Normal comparison is either a **cmp** instruction or an instruction with conditional modifier
- Special comparison **cmpn** with conditional modifier of EQ, GT, GE, LT, LE, when the second source operand is NaN, returns TRUE, regardless of the first source operand, and when the second source operand is not NaN, but first one is, returns FALSE. **Cmpn** of NE, when the second source operand is NaN, returns FALSE, regardless of the first source operand, and when the second source operand is not NaN, but first one is, returns TRUE.
  - This is used to support the proposed IEEE-754R rule on **min** or **max** operations. For which, if only one operand is NaN, min and max operations return the other operand as the result.
- Both normal and special comparisons of any non-NaN value against +/- INF return exact result according to the conditional modifier. This is because that infinities are exact representation in the sense that +INF = +INF and -INF = -INF.
  - NaN is unordered in the sense that NaN != NaN.
- IEEE-754 requires floating point operations to produce a result that is the nearest representable value to an infinitely precise result, known as "round to nearest even" (RTNE). 32-bit floating point operations must produce a result that is within 0.5 Unit-Last-Place (0.5 ULP) of the infinitely precise result. This applies to addition, subtraction, and multiplication.
- All arithmetic floating point instructions does Round To Nearest Even at the end of the computation, except the round instructions.

#### 4.3.1.2 Complete Listing of Deviations or Additional Requirements vs. IEEE-754

For a result that cannot be represented precisely by the floating point format, the execution unit uses rounding to nearest even to produce a result to an infinitely precise result. This ends up with a result that is within 0.5 Unit-Last-Place( 0.5 ULP) of the infinitely precise result.

- **[DevSNB]** Denorms are flushed to sign-preserved zero on input and output of any floating point mathematical operation.
  - The exception to the above point about flushing denorms is that any I/O or data movement operation that does not manipulate the data (such as point sampling float data, or executing any raw "mov" instruction, or any sort of conditional raw "mov" if present) must not alter data at all (so a denorm remains denorm). Note that doing something that amounts to just moving data, but isn't explicitly, such as multiplying a number by 1.0f is not detected as just a raw "mov", and in this case a denorm flush would happen.



- NaN input to an operation obviously always produces NaN on output, however the exact bit pattern of the NaN is not required to stay the same (unless the operation is a raw “mov” instruction which does not alter data at all.)
- $x * 1.0f$  must always result in  $x$  (except denorm flushed and possible bit pattern change for NaN).
- $x +/- 0.0f$  must always result in  $x$  (except denorm flushed and possible bit pattern change for NaN). But  $-0 + 0 = +0$ .
- Fused operations (such as mac, dp4, dp3, etc.) may produce intermediate results out of 32-bit float range, but whose final results would be within 32-bit float range if intermediate results were kept at greater precision. In this case, implementations are permitted to produce either the correct result, or else +/-INF. Thus, compatibility between a fused operation, such as “mac”, with the unfused equivalent, “mul” followed by “add” in this case, is not guaranteed.
  - As the accumulator registers have more precision than 32-bit float, any instruction with accumulator as a source/destination operand may produce a different result than that using GRF/MRF registers.
- API Shader divide operations are implemented as  $x * (1.0f/y)$ . With the two-step method,  $x * (1.0f/y)$ , the multiply and the divide each independently operate at the 32-bit floating point precision level (accuracy to 1 ULP).
- See the Type Conversion section below for rules on converting to/from float representations.



### 4.3.1.3 Comparison of Floating Point Numbers

The following tables (Table 4-3 through Table 4-8) detail the rules for floating point comparison. In the tables, “+/-Fin” stands for a positive or negative finite precision floating point number. Result is either a true (T) or false (FALSE or F). Each row corresponds to a fixed <src0> and each column corresponds to a fixed <src1>. When comparing two positive finite numbers (or two negative finite numbers), the result can be T or F depending on the values. Therefore, the corresponding fields in the following tables are marked as T/F.

**Table 4-3. Results of “Greater-Than” Comparison – CMP.G**

src0	src1	-inf	-Fin	-denorm	-0	+0	+denorm	+Fin	+inf	NaN
-inf		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
-Fin		T	T/F	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
-denorm		T	T	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
-0		T	T	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
+0		T	T	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
+denorm		T	T	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
+Fin		T	T	T	T	T	T	T/F	FALSE	FALSE
+inf		T	T	T	T	T	T	T	FALSE	FALSE
NaN		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE

**Table 4-4. Results of “Less-Than” Comparison – CMP.L**

src0	src1	-inf	-Fin	-denorm	-0	+0	+denorm	+Fin	+inf	NaN
-inf		FALSE	T	T	T	T	T	T	T	FALSE
-Fin		FALSE	T/F	T	T	T	T	T	T	FALSE
-denorm		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T	T	FALSE
-0		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T	T	FALSE
+0		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T	T	FALSE
+denorm		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T	T	FALSE
+Fin		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T/F	T	FALSE
+inf		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
NaN		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE



**Table 4-5. Results of “Equal-To” Comparison – CMP.E**

src0	src1	-inf	-Fin	-denorm	-0	+0	+denorm	+Fin	+inf	NaN
-inf		T	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
-Fin		FALSE	T/F	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
-denorm		FALSE	FALSE	T	T	T	T	FALSE	FALSE	FALSE
-0		FALSE	FALSE	T	T	T	T	FALSE	FALSE	FALSE
+0		FALSE	FALSE	T	T	T	T	FALSE	FALSE	FALSE
+denorm		FALSE	FALSE	T	T	T	T	FALSE	FALSE	FALSE
+Fin		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T/F	FALSE	FALSE
+inf		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T	FALSE
NaN		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE

**Table 4-6. Results of “Not-Equal-To” Comparison – CMP.NE**

src0	src1	-inf	-Fin	-denorm	-0	+0	+denorm	+Fin	+inf	NaN
-inf		FALSE	T	T	T	T	T	T	T	T
-Fin		T	T/F	T	T	T	T	T	T	T
-denorm		T	T	FALSE	FALSE	FALSE	FALSE	T	T	T
-0		T	T	FALSE	FALSE	FALSE	FALSE	T	T	T
+0		T	T	FALSE	FALSE	FALSE	FALSE	T	T	T
+denorm		T	T	FALSE	FALSE	FALSE	FALSE	T	T	T
+Fin		T	T	T	T	T	T	T/F	T	T
+inf		T	T	T	T	T	T	T	FALSE	T
NaN		T	T	T	T	T	T	T	T	T

**Table 4-7. Results of “Less-Than Or Equal-To” Comparison – CMP.LE**

src0	src1	-inf	-Fin	-denorm	-0	+0	+denorm	+Fin	+inf	NaN
-inf		T	T	T	T	T	T	T	T	FALSE
-Fin		FALSE	T/F	T	T	T	T	T	T	FALSE
-denorm		FALSE	FALSE	T	T	T	T	T	T	FALSE
-0		FALSE	FALSE	T	T	T	T	T	T	FALSE
+0		FALSE	FALSE	T	T	T	T	T	T	FALSE
+denorm		FALSE	FALSE	T	T	T	T	T	T	FALSE
+Fin		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T/F	T	FALSE
+inf		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T	FALSE
NaN		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE



**Table 4-8. Results of “Greater-Than or Equal-To” Comparison – CMP.GE**

src0	src1	-inf	-Fin	-denorm	-0	+0	+denorm	+Fin	+inf	NaN
-inf		T	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
-Fin		T	T/F	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
-denorm		T	T	T	T	T	T	FALSE	FALSE	FALSE
-0		T	T	T	T	T	T	FALSE	FALSE	FALSE
+0		T	T	T	T	T	T	FALSE	FALSE	FALSE
+denorm		T	T	T	T	T	T	FALSE	FALSE	FALSE
+Fin		T	T	T	T	T	T	T/F	FALSE	FALSE
+inf		T	T	T	T	T	T	T	T	FALSE
NaN		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE

#### 4.3.1.4 Min/Max of Floating Point Numbers

A special comparison called Compare-NaN is introduced to handle the difference of above mentioned floating point comparison and the rules on supporting MIN/MAX. To compute the MIN or MAX of two floating point numbers, if one of the numbers is NaN and the other one is not, MIN or MAX of the two numbers returns the one that is not NaN. When two numbers are NaN, MIN or MAX of the two numbers always returns source1.

Mix and Max is achieved using conditional selects, i.e., SEL with ‘condition modifiers’ as :

Evaluations	Instructions
$\text{MIN}(\text{src0}, \text{src1}) = (\text{src0} < \text{src1}) ? \text{src0} : \text{src1}$	sel.l.f0.0 dst src0 src1
$\text{MAX}(\text{src0}, \text{src1}) = (\text{src0} \geq \text{src1}) ? \text{src0} : \text{src1}$	sel.ge.f0.0 dst src0 src1

Note even f0.0 is specified in the instruction, the flag register is not touched by this instruction.

The following tables (**Error! Reference source not found.** through Table 4-10) detail the rules for this special compare-NaN operation for floating point numbers. Notice that excepting “Not-Equal-To” comparison-NaN, last columns in all other tables have ‘T’.



**Table 4-9. Results of “Less-Than” Comparison-NaN – CMPN.L**

src0	src1	-inf	-Fin	-denorm	-0	+0	+denorm	+Fin	+inf	NaN
-inf		FALSE	T	T	T	T	T	T	T	T
-Fin		FALSE	T/F	T	T	T	T	T	T	T
-denorm		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T	T	T
-0		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T	T	T
+0		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T	T	T
+denorm		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T	T	T
+Fin		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T/F	T	T
+inf		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T
NaN		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE

**Table 4-10. Results of “Greater-Than or Equal-To” Comparison-NaN – CMPN.**

src0	src1	-inf	-Fin	-denorm	-0	+0	+denorm	+Fin	+inf	NaN
-inf		T	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T
-Fin		T	T/F	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T
-denorm		T	T	T	T	T	T	FALSE	FALSE	T
-0		T	T	T	T	T	T	FALSE	FALSE	T
+0		T	T	T	T	T	T	FALSE	FALSE	T
+denorm		T	T	T	T	T	T	FALSE	FALSE	T
+Fin		T	T	T	T	T	T	T/F	FALSE	T
+inf		T	T	T	T	T	T	T	T	T
NaN		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE

**Table 4-11. Results of “Greater-Than or Equal-To” Comparison-NaN – CMPN.GE**

src0	src1	-inf	-Fin	-denorm	-0	+0	+denorm	+Fin	+inf	NaN
-inf		T	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T
-Fin		T	T/F	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	T
-denorm		T	T	T	T	T	T	FALSE	FALSE	T
-0		T	T	T	T	FALSE	T	FALSE	FALSE	T
+0		T	T	T	T	T	T	FALSE	FALSE	T
+denorm		T	T	T	T	T	T	FALSE	FALSE	T
+Fin		T	T	T	T	T	T	T/F	FALSE	T
+inf		T	T	T	T	T	T	T	T	T
NaN		FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE





## 4.3.2 Alternative Floating Point Mode

The key characteristics of the alternative floating point mode is that NaN, Inf and denorm are not expected for an application to pass into the graphics pipeline, and the graphics hardware must not generate NaN, Inf or denorm as computation result. For example, a result that is larger than the maximum representable floating point number is expected to be flushed to the largest representable floating point number, i.e., +FLT\_MAX. The FLT\_MAX has an exponent of 0xFE and a mantissa of all one's, which is the same for IEEE floating point mode.

Here is the complete list of the differences of legacy graphics mode from the relaxed IEEE-754 floating point mode.

- Any +/- INF result must be flushed to +/- FLT\_MAX, instead of being output as +/- INF.
- Extended mathematics functions of log(), rsq() and sqrt() take the absolute value of the sources before computation to avoid generating INF and NaN results.

Table 4-12 shows the support of these differences in various hardware units.

**Table 4-12. Supported Legacy Float Mode and Impacted Units**

IEEE-754 Deviations	VF	Clipper	SF	WIZ	EU	EM	Sampler	RC
Any +/- INF result flushed to +/- FLT_MAX	Y	Y	Y	Y	Y	Y	Y	Y
Log, rsq, sqrt take abs() of sources	N/A	N/A	N/A	N/A	N/A	Y	N/A	N/A

Table 4-13 shows some of the desired or recommended alternative floating point mode behaviors that do not have hardware design impact. The reasons of not needing special hardware support for these items are also provided.



**Table 4-13. Dismissed legacy behaviors**

<b>Suggested IEEE-754 Deviations</b>	<b>Reason for Dismiss</b>
Mov forces (+/-)INF to (+/-)FLT_MAX	(+/-)INF is never present as input
(+/-)INF – (+/-)INF = +/- FLT_MAX instead of NaN	(+/-)INF is never present as input
Denorm must be flushed to zero in all cases (including trivial mov and point sampling)	Denorm is never present as input
Anything*0=0 (including NaN*0=0 and INF*0=0)	NaN and INF are never present as input
Except propagated NaN, NaN is never generated	NaN is never present as input and the system never generates NaN based on rules in the previous table
An input NaN gets propagated excepting (a)-(d)	NaN is never present as input
(a) Rcp (and rsq) of 0 yields FLT_MAX	N/A, as it is already covered by the general rule “Any +/- INF result flushed to +/- FLT_MAX”
(b) Sampler honors 0/0 = 0 as if (1/0)*0	There is no divide in Sampler
l Rcp (and rsq) of INF yields +/- 0	(+/-)INF is never present as input
(d) Sampler honors INF/INF = 0 as if (1/INF)=0 followed by Anything*0 = 0	There is no divide in Sampler

## 4.4 Type Conversion

### 4.4.1 Float to Integer

Converting from float to integer is based on rounding toward zero. If the floating point value is +0, -0, +Denorm, -Denorm, +NaN –r -NaN, the resulting integer value is always 0. If the floating point value is positive infinity (or negative infinity), the conversion result takes the largest (or the smallest) represent-able integer value. If the floating point value is larger (or smaller) than the largest (or the smallest) represent-able integer value, the conversion result takes the largest (or the smallest) represent-able integer value. The following table shows these special cases. The last two rows are just examples. They can be any number outside the represent-able range of the output integer type (UD, D, UW, W, UB and B).

<b>Input Format</b>	<b>Output Format</b>					
<b>F</b>	<b>UD</b>	<b>D</b>	<b>UW</b>	<b>W</b>	<b>UB</b>	<b>B</b>
+/- Zero	00000000	00000000	00000000	00000000	00000000	00000000
+/- Denorm	00000000	00000000	00000000	00000000	00000000	00000000
NAN	00000000	00000000	00000000	00000000	00000000	00000000
-NAN	00000000	00000000	00000000	00000000	00000000	00000000
INF	FFFFFFFF	7FFFFFFFF	0000FFFF	00007FFF	000000FF	0000007F



Input Format	Output Format					
	00000000	80000000	00000000	00008000	00000000	00000080
-INF	00000000	80000000	00000000	00008000	00000000	00000080
$+2^{32}$ (*)	FFFFFFFF	7FFFFFFFF	0000FFFF	00007FFF	000000FF	0000007F
$-2^{32}-1$ (*)	00000000	80000000	00000000	00008000	00000000	00000080

## 4.4.2 Integer to Integer with Same or Higher Precision

Converting an unsigned integer to a signed or an unsigned integer with higher precision is based on zero extension.

Converting an unsigned integer to a signed integer with the same precision is based on modular wrap-around. Without saturation, a larger than represent-able number becomes a negative number. With saturation, a larger than represent-able number is saturated to the largest positive represent-able number.

Converting a signed integer to a signed integer with higher precision is based on sign extension.

Converting a signed integer to an unsigned integer with higher precision is based on zero extension. Without saturation, a negative number becomes a large positive number with the sign bit wrapped-up. With saturation, a negative number is saturated to zero.

## 4.4.3 Integer to Integer with Lower Precision

Converting a signed or an unsigned integer to a signed or an unsigned integer with lower precision is based on bit truncation. Without saturation, only the lower bits are kept in the output regardless of the sign-ness of input and output. With saturation, a number that is outside the represent-able range is saturated to the closest represent-able value.

## 4.4.4 Integer to Float

Converting a signed or an unsigned integer to a single precision float number is to round to the closest representable float number. For any integer number with magnitude less than or equal to 24 bits, resulting float number is a precise representation of the input. However, if it is more than 24 bits, by default a “round to nearest even” is performed.



## 5. Execution Environment

### 5.1 Overview

The instruction set is a general-purpose data-parallel instruction set optimized for graphics and media computations. Supports for 3D graphics API (application program interface) Shader instructions are mostly native, meaning that DevSNB provides efficient execution for Shader programs. Depending on the Shader program operation modes (for example, a Vertex Shader may be executed on a base of a vertex-pair, while a Pixel Shader may be executed on a base of a 16-pixel group), translation from 3D graphics API Shader instruction streams into native instructions may be required. In addition, there are many specific capabilities to accelerate media applications. The following list provides a summary of the instruction set.

- ISA support SIMD (single instruction multiple data) instructions. The number of data elements per instruction depends on the data type.
- ISA supports SIMD parallel arithmetic, vector arithmetic, logical, and SIMD control/branch instructions.
- ISA supports instruction level variable-width SIMD execution.
- ISA supports conditional SIMD execution via destination mask, predication, and execution mask.
- ISA supports instruction compaction.
- An instruction may be executed in multiple cycles over a SIMD execution pipeline.
- Most instructions have three operands. Some instructions have additional implied source and destination operands. Some instructions have explicit dual destinations.
- ISA supports region-based register addressing.
- ISA supports direct and indirect (indexed) register addressing.
- instructions may have a scalar and vector immediate source operand.
- Higher precision accumulator registers are architecturally visible.
- Self-modifying code is not allowed (instruction streams, including instruction caches, are read-only).

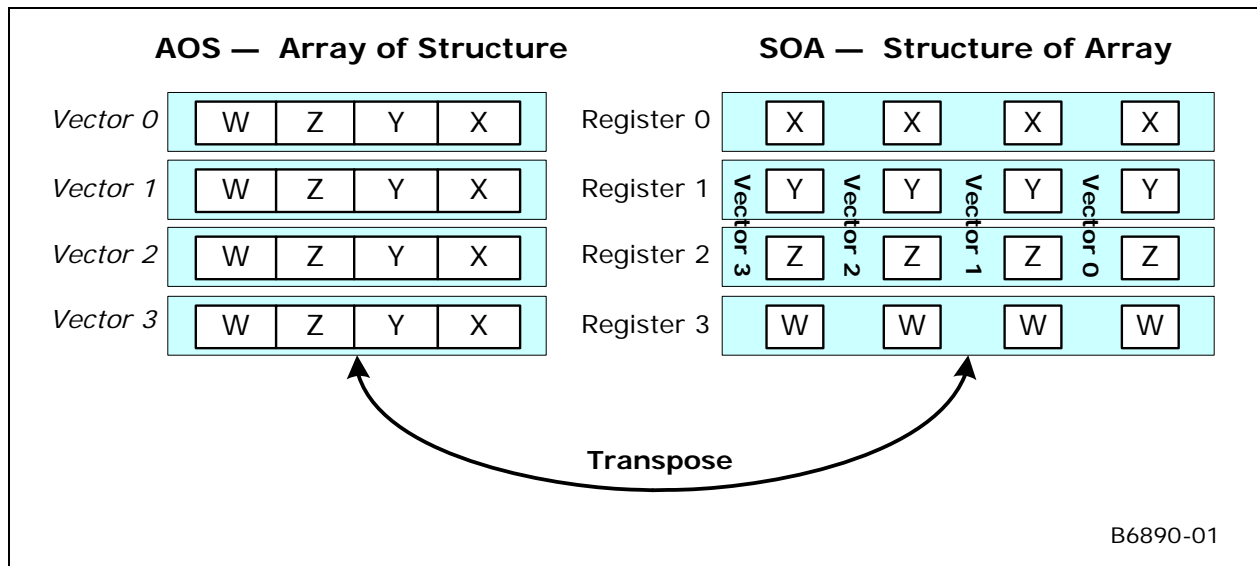
## 5.2 Primary Usage Models

In describing the usage models of the instruction set, it is inevitable to forward reference terminology, syntax and instructions detailed later in this specification. For clarity reasons, not all forward references will be provided in this section as well as subsequent sections. For example, reference to binary instruction fields such as **Align1**, **Align16**, **Compr**, **SecHalf**, etc, can be found in the Instruction Summary chapter. And assembly instruction syntax can be found in the Instruction Summary chapter and Instruction Reference chapter.

### 5.2.1 AOS and SOA Data Structures

With the Align1 and Align16 access modes, the instruction set provides effective SIMD computation regardless whether data are arranged in array of structure (AOS) form or in structure of array (SOA) form. The AOS and SOA data structures are illustrated by the examples in Figure 5-1. The example shows two different ways of storing four vectors in four SIMD registers. For simplicity, data vector and SIMD register both have four data elements. The four data elements in a vector are denoted by X, Y, Z and W just as for a vertex in 3D geometry. The AOS structure stores one vector in a register and the next vector in another register. The SOA structure stores one data element of each vector in a register and the next element of each vector in the next register and so on. It is obvious in this case the two structures can be related by a matrix transpose operation.

Figure 5-1. AOS and SOA data structures



3D and media applications take advantage of such broad architecture support and use both AOS and SOA data arrangements.

- Vertices in 3D Geometry (Vertex Shader and Geometry Shader) are arranged in AOS structure and run on SIMD4x2 and SIMD4 modes, respectively, as detailed below.
- Pixels in 3D Rasterization (Pixel Shader) are arranged in SOA structure and run on SIMD8 and SIMD16 modes as detailed below.



- Pixels in media are primarily arranged in SOA structure, and occasionally in AOS structure with possible mixed mode of operations that use region-based addressing extensively.

These are preferred methods; alternative arrangements may also be possible. Shared function resources provide data transpose capability to support both modes of operations: The sampler has a transpose for sample reads, the data port has a transpose for render cache writes, and the URB unit has a transpose for URB writes.

The following 3D graphics API Shader instruction will be used in the following sections to illustrate various modes of operations:

```
add <dst>.xyz <src0>.yxzw <src1>.zwxy
```

This example is an SIMD instruction that takes two source operands `<src0>` and `<src1>`, performs addition operation (add), and store the additions to the destination operand `<dst>`. Each operand contains four floating point data elements. The data type is determined by the instruction opcode. This instruction also uses source swizzle modifier `.yxzw` for `<src0>` and `.zwxy` for `<src1>` and destination mask modifier `.xyz`. Please refer to programming specifications of 3D graphics API Shader instructions for more details.

A physical GRF register has 256 bits, which may be used to store 8 floating point data elements. For 3D graphics usage, the mode of operation is (loosely) termed after the data structure as `SIMD $m$ x $n$` , where “ $m$ ” is a numerical term describing the size of vector and “ $n$ ” is the number of concurrent program flows executed in SIMD.

- Execution with AOS data structures
  - **SIMD4 (short for SIMD4x1)** stands for the mode of operation where a SIMD instruction operates on 4-element vectors stored packed in the registers. There is only one program flow.
  - **SIMD4x2** standards for the SIMD operation based on a pair of 4-element vectors stored in a register. There are effectively two programs running side by side with one vector per program.
- Execution with SOA data structures – also referred to as “channel serial” execution
  - **SIMD8 (short for SIMD1x8)** standards for the SIMD operation based on the SOA data structure where one register contains one data element (the same one) of 8 vectors. Effectively, there are 8 concurrent program flows.
  - **SIMD16 (short for SIMD1x16)** is a special term indicating the use of instruction compression whereas each compressed SIMD instruction operates on a pair of registers that contains one data element (the same one) of 16 vectors. SIMD16 has 16 concurrent program flows.

## 5.2.2 SIMD4 Mode of Operation

With a register mapping of `<src0>` to doublewords 0-3 of `r2`, `<src1>` to doublewords 4-7 of `r2` and `<dst>` to doublewords 0-3 of `r3`, the example 3D graphics API Shader instruction can be translated into the following instruction:

```
add (4) r3<4>.xyz:f r2<4>.yzwx:f r2.4<4>.zwxy:f {NoMask}
```


Without diving too much into the syntax definition of an instruction, it is clear that an instruction also takes two source operands and one destination operands. The second term, (4), is the execution size that determines the number of data elements processed by the SIMD instruction. It is similar to the term SIMD Width used in the literature. Each operand is described by the register region parameters such as '<4>' and data type (e.g. ":f"). These will be detailed in Section 5.3. The instruction option field, {NoMask}, ensure that the execution occurs for the execution channels shown in the instruction, instead of, possibly, being masked out by the conditional masks of the thread (See Instruction Summary chapter for definition of *MaskCtrl* instruction field).

The operation of this instruction is illustrated in Figure 5-2. In this example, both source operands share the same physical GRF register r2. The two are distinguished by the subregister number. The source swizzles control the routing of source data elements to the parallel adders corresponding to the destination data elements. The shaded areas in the destination register r3 are not modified. In particular, doublewords 4-7 are unchanged as the execution size is 4; doubleword 3 is unchanged due to the destination mask setting.

In this mode of operation, there is only one program flow – any branch decision will be based on a scalar condition and apply to the whole vector of four elements. Option {NoMask} ensures that the instruction is not subject to the masks. In fact, most of the instructions in a thread should have {NoMask} set.

Even though the execution only performs four parallel add operations, the GEN instruction still executes in 2 cycles (with no useful computation in the second cycle).

**Figure 5-2. A SIMD4 Example**



### 5.2.3 SIMD4x2 Mode of Operation

In this mode, two corresponding vectors from the two program flows fill a physical register. With a register mapping of <src0> to r2, <src1> to r3 and <dst> to r4, the example 3D graphics API Shader instruction can be translated into the following instruction:

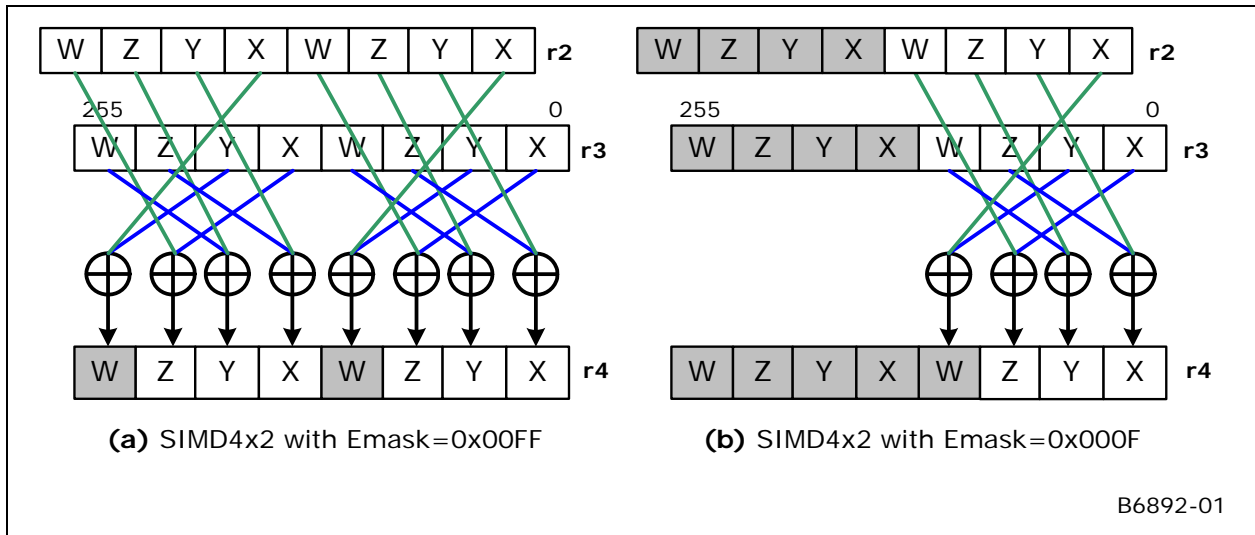
```
add (8) r4<4>.xyz:f r2<4>.yxzw:f r3<4>.zwxy:f
```

This instruction is subject to the execution mask, which initiated from the dispatch mask. If both program flows are available (e.g. Vertex Shader executed with two active vertices), the dispatch mask is set to 0x00FF. The operation of this instruction is illustrated in Figure 5-3 (a). The source swizzles control the routing of source data elements to the parallel adders corresponding to the destination data elements. The shaded areas in the destination register r3 (doublewords 3 and 7) are unchanged due to the



destination mask setting. If only one program flow is available (e.g. the same SIMD4x2 Vertex Shader with only one active vertex), the dispatch mask is set to 0x000F. The operation of the same instruction is shown in Figure 5-3(b).

**Figure 5-3. SIMD4x2 Examples with Different Emasks**




The two source operands only need to be 16-byte aligned, not have to be GRF register aligned. For example, the first source operand could be a 4-element vector (e.g. a constant) stored in doublewords 0-3 in r2, which is shared by the two program flows. The example 3D graphics API Shader instruction can then be translated into the following instruction:

```
add (8) r4<4>.xyz:f r2<0>.yzwx:f r3<4>.zwxxy:f
```

The only difference here is that the vertical stride of the first source is 0. The operation of this instruction is illustrated in Figure 5-4.

**Figure 5-4. A SIMD4x2 Example with a Constant Vector Shared by Two Program Flows**





## 5.2.4 SIMD16 Mode of Operation

With 16 concurrent program flows, one element of a vector would take two GRF registers. In this mode, two corresponding vectors from the two program flows fill a physical register.

With the following register mappings,

<src0>:            r2-r9 (with 16 X data elements in r2-r3, Y in r4-5, Z in r6-7 and W in r8-9),

<src1>:            r10-r17,

<dst>:             r18-r25,

the example 3D graphics API Shader instruction can be translated into the following three instructions:

`add (16) r18<1>:f r4<8;8,1>:f r14<8;8,1>:f            // dst.x = src0.y + src1.z`

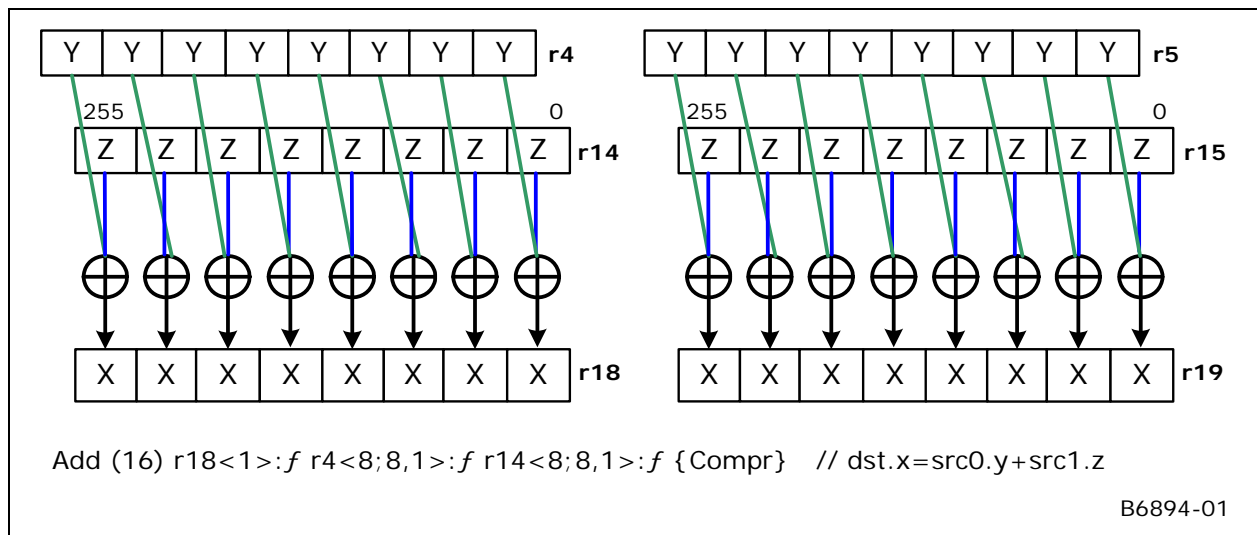
`add (16) r20<1>:f r6<8;8,1>:f r16<8;8,1>:f            // dst.y = src0.z + src1.w`

`add (16) r22<1>:f r8<8;8,1>:f r10<8;8,1>:f            // dst.z = src0.w + src1.x`

The three instructions correspond to the three enabled destination masks. As there is no output for the W elements of <dst>, no instruction is needed for that element. The first instruction inputs the Y elements of <src0> and the Z elements of <src1> and outputs the X elements of <dst>. The operation of this instruction is shown in Figure 5-5.

With the number of program flows more than one, the above instructions also subject to execution mask. The 16-bit dispatch mask is partitioned into four groups with four bits each. For Pixel Shader generated by the Windower, each 4-bit group corresponds to a 2x2 pixel subspan. If a subspan is not valid for a Pixel Shader instance, the corresponding 4-bit group in the dispatch mask is not set. Therefore, the same instructions can be used independent of the number of available subspans without creating bogus data in the subspans that are not valid.

**Figure 5-5. A SIMD16 Example**






Similar to SIMD4x2 mode, a constant may also be shared for the 16 program flows. For example, the first source operand could be a 4-element vector (e.g. a constant) stored in doublewords 0-3 in r2 (AOS format). The example 3D graphics API Shader instruction can then be translated into the following instruction:

```

add (16) r18<1>:f r2.1<0;1,0>:f r14<8;8,1>:f {Compr} // dst.x = src0.y + src1.z
add (16) r20<1>:f r2.2<0;1,0>:f r16<8;8,1>:f {Compr} // dst.y = src0.z + src1.w
add (16) r22<1>:f r2.3<0;1,0>:f r10<8;8,1>:f {Compr} // dst.z = src0.w + src1.x
  
```

The register region of the first source operand represents a replicated scalar. The operation of the first instruction is illustrated in Figure 5-6.

**Figure 5-6. Another SIMD16 Example with an AOS Shared Constant**



```

Add (16) r18<1>:f r2.1<0;1,0>:f r14<8;8,1>:f {Compr} // dst.x=src0.y+src1.z
  
```

B6895-01

### 5.2.5 SIMD8 Mode of Operation

Each compressed instruction has two correspond uncompressed instructions. Taking the example instruction shown in Figure 5-6, it is equivalent to the following two instructions.

```

add (8) r18<1>:f r4<8;8,1>:f r14<8;8,1>:f // dst.x[7:0] = src0.y + src1.z
add (8) r19<1>:f r5<8;8,1>:f r15<8;8,1>:f {SecHalf} // dst.x[15:8] = src0.y + src1.z
  
```

Therefore, SIMD8 can be viewed as a special case for SIMD16.

There are other reasons that SIMD8 instructions may be used. Within a program with 16 concurrent program flows, some time SIMD8 instruction must be used due to architecture restrictions. For example, the address register a0 only have 8 elements, if an indirect GRF addressing is used, SIMD16 instructions are not allowed.



## 5.3 Registers and Register Regions

### 5.3.1 Register Files

Registers are grouped into different name spaces called register files. There are three different register files defined: General Register File, Message Register File, and Architecture Register File. In addition, immediate operands also have a unique encoding of the register file field, even though they come inline in the instruction word and do not have dedicated physical storages.

- General Register File (GRF): GRF contains general-purpose read-write registers.
- Message Register File (MRF): MRF contains special purpose registers used for message passing only. MRF registers are write-only.
- Architecture Register File (ARF): ARF contains all other architectural registers, including the address registers (*a#*), accumulators (*acc#*), flags (*f#*), masks (*mask#*), mask stack (*ms#*), mask stack depth (*msd#*), notification count (*n#*), instruction pointer (*ip*), etc. Null register (*null*) is also encoded as an ARF register.
- Immediate: Certain instructions take immediate terms as the source operands. These immediate terms have a distinct register file encoding.

Each thread executed in an EU has its own thread context, i.e. dedicated register space, which is not shared between threads executing on a common EU or on a different EU. In the rest of the Chapters, register access are in respect to a given thread.

### 5.3.2 GRF Registers

Number of Registers:	Various
Default Value:	None
Normal Access:	RW
Elements:	Various
Element Size:	Various
Element Type:	Various
Access Granularity:	Byte
Write Mask Granularity:	Byte
Index-ability:	Yes

Registers in the General Register File are the most commonly used read-write registers. During the execution of a thread, GRF registers are used to store the temporary data, and serve as the destination to receive data from shared function units (and some times from a fixed function unit). They are also used to store the input (initialization) data when a thread is created. By allowing fixed function hardware to initialize some portion of GRF registers during thread dispatch time, the architecture can achieve better parallelism. A thread's execution efficiency can also be improved as some data are already in the register to be executed upon. Besides these registers containing thread's payload, the rest of GRF registers of a thread are not initialized.



**Table 5-1. Summary of GRF Registers**

<b>Register File</b>	<b>Register Name</b>	<b>Description</b>
<i>General Register File (GRF)</i>	<b><i>r#</i></b>	General purpose read write registers

Each execution unit has a fixed size physical GRF register RAM. The GRF register RAM is shared by all threads on the EU. GRF space for a thread is allocated at thread dispatch time, allowing the amount of GRF space to adapt to the need of a given thread.

Mapping of a thread's GRF registers to the physical GRF RAM is through a translation table. Therefore, a thread's access to GRF is always through the 0-based logical view. For example, the GRF registers are *r0* through *r127*.

GRF registers can be accessed using region-based addressing at byte granularity (both read and write). A source operand must be contained within two adjacent physical registers. A destination operand must be contained within one physical register. GRF registers support direct addressing and register-indirect addressing. Register-indirect addressing uses the address registers (ARF registers *a#*) and an immediate address offset value.

When accessing (read and/or write) outside the GRF register range allocated for a given thread either through direct or indirect addressing, the result is unpredictable.

**Table 5-2. GRF Registers Available in Device Hardware**

<b>Device</b>	<b>Physical Register Size</b>	<b>Allocation Granularity</b>	<b>Number per Thread</b>	<b>Number per EU</b>
[DevSNB]	256 bits	Fixed allocation of 128 register	128 registers	640 registers



### 5.3.3 MRF Registers

Number of Registers:	Fixed
Default Value:	None
Normal Access:	WO
Elements:	Various
Element Size:	Various
Element Type:	Various
Access Granularity:	Byte
Write Mask Granularity:	Byte
Index-ability:	Yes

Registers in the Message Register File are used to store the header and payload for out-going messages from a thread to a shared function such as the Sampler. There are fixed number of MRF registers for each thread.

MRF registers are write-only, and therefore, can only be the destination operand of an instruction.

MRF registers support write-enable at byte granularity. When an MRF register is used as the current destination operand of the *send* instruction, only 256-bit register aligned access is supported.

When accessing (write) outside the MRF register range for a given thread, the result is unpredictable.

**Table 5-3. Summary of MRF Registers**

Register File	Register Name	Description
Message Register File (MRF)	<i>m#</i>	Special purpose output write-only registers

**Table 5-4. MRF Registers Available in Device Hardware**

Device	Physical Register Size	Number per Thread	Indirect Addressing?
[DevSNB]	256 bits	24 registers	Yes

### 5.3.4 ARF Registers

#### 5.3.4.1 Overview

Besides GRF and MRF registers that are directly indicated by unique register file coding, all other registers belong to the general Architecture Register File (ARF). Encoding of architecture register types are based on the MSBs of the register number field, RegNum, in the instruction word. RegNum field has 8 bits. The 4 MSBs, RegNum[7:4], represent the architecture register type.



**Table 5-5. Summary of Architecture Registers [DevSNB+]**

Register Type (RegNum [7:4])	Register Name	Register Count	Description
0000	<i>null</i>	1	Null register
0001	<i>a0.#</i>	1	Address register
0010	<i>acc#</i>	2	Accumulator register
0011	<i>f0.#</i>	1	Flag register [DevSNB]
0011	<i>reserved</i>		
0100	<i>reserved</i>		
0110	<i>reserved</i>		
0111	<i>sr0.#</i>	1	State register
1000	<i>cr0.#</i>	1	Control register
1001	<i>n#</i>	2	Notification count register
1010	<i>ip</i>	1	Instruction pointer register
1011	<i>tdr</i>	1	Thread dependency register
1100	<i>reserved</i>		
1101	<i>reserved</i>		
1111	<i>reserved</i>		

The remaining register number field RegNum[3:0] is used to identify the register number of a given architecture register type. Therefore, maximum number of registers for a given architecture register type is limited to 16. The subregister number field, SubRegNum, in instruction word has 5 bits. It is used for addressing subregister region for an architecture register supporting register subdivision. SubRegNum field is in unit of byte. Therefore, maximum number of bytes of an architecture register is limited to 32. Depending on alignment restriction of a register type, only certain encodings of SubRegNum field is applicable for an architecture register. The detailed definitions are provided in the following sections.

### 5.3.4.2 Access Granularity

ARF registers may be accessed with subregister granularity according to the descriptions below and following the same rule of region-based addressing for GRF and MRF. The machine code for register number and subregister number of ARF follows the same rule as for other register files with byte granularity. For an ARF as a source operand, the region-based address controls the source swizzle mux. The destination subregister number and destination horizontal stride can be used to control to generate the destination write mask at byte level.

A special restriction on region-based addressing for ARF is that the register region cannot cross register boundary. This rule in fact only applies to the accumulator as it is the only ARF register containing multiple registers (two).

Subregister fields of an ARF register may not all populated (indicated by the subregister indicated as reserved). Write to an unpopulated subregister will be dropped, there is no side effect. Read from an unpopulated subregister, if not specified, will return unpredictable data.



Some of ARF registers are read-only. Write to a read-only ARF register is dropped and there is no side effect.

### 5.3.4.3 Null Register

ARF Register Type Encoding (RegNum[7:4]):	0000b
Number of Registers:	1
Default Value:	N/A
Normal Access:	N/A
Elements:	N/A
Element Size:	N/A
Element Type:	N/A
Access Granularity:	N/A
Write Mask Granularity:	N/A
SecHalf Control:	N/A
Index-ability:	No

The null register is a special encoding for an operand that does not have physical map. It is primarily used in the instruction to indicate the non-existence of an operand. Write to the null register has no side effect. Read from the null register returns undefined result.

The null register can be used in the place when a source operand is absent. For example, for a single source operand instruction such as MOV, NOT, the second source operand <src1> must be a null register.

When the null register is used as the destination operand of an instruction, it indicates the computed results are not stored in any physical registers. However, implied writes to the accumulator register, if applicable, may still occur for the instruction. When the conditional modifier is present, update to the selected flag register also happens. In this case, the register region fields of the 'null' operand are valid.

Another example use is to use the null register as the posted destination of a *send* instruction for data write to indicate that there is no write completion acknowledgement required. In this case, however, the register region fields are still valid. The null register can also be the first source operand for a *send* instruction indicating the absent of the implied move. See *send* instruction for details.



### 5.3.4.4 Address Register

ARF Register Type Encoding (RegNum[7:4]):	0001b
Number of Registers:	1
Default Value:	None
Normal Access:	RW
Elements:	8
Element Size:	16 bits
Element Type:	UW or UD
Access Granularity:	Word
Write Mask Granularity:	Word
SecHalf Control:	N/A
Index-ability:	No

There are eight address subregisters forming an 8-element vector. Each address subregister contains 16 bits. Address subregisters can be used as regular source and destination operands, as the indexing addresses for register-indirect-addressed access of GRF registers, and also as the source of the message descriptor for the *send* instruction.

**Table 5-6. Register and Subregister Numbers for Address Register**

RegNum[3:0]	SubRegNum[4:0]
0000 = <b>a0</b> All other encodings are reserved.	<p>When register a0 or subregisters in a0 is used as the address register for register-indirect addressing, the address subregisters must be accessed as unsigned word integers. Therefore, the subregister number field must also be word-aligned.</p> <p>00000 = <b>a0.0:uw</b>            00010 = <b>a0.1:uw</b>            00100 = <b>a0.2:uw</b>            00110 = <b>a0.3:uw</b>            01000 = <b>a0.4:uw</b>            01010 = <b>a0.5:uw</b>            01100 = <b>a0.6:uw</b>            01110 = <b>a0.7:uw</b></p> <p>All other encodings are reserved.</p> <p>However, when register a0 or subregisters in a0 is an explicit source and/or destination register, other data types are allowed as long as the register region falls in the 128-bit range.</p>





Table 5-7. Address Register Fields

Dword	Bits	Subfield Description
3	31:16	<b>Address subregister a0.7:uw.</b> This field, with only the lower 12 bits populated, can be used as an unsigned integer for register-indirect register addressing. Format: U12
	15:0	<b>Address subregister a0.6:uw.</b> This field, with only the lower 12 bits populated, can be used as an unsigned integer for register-indirect register addressing. Format: U12
2	31:16	<b>Address subregister a0.5:uw.</b> This field, with only the lower 12 bits populated, can be used as an unsigned integer for register-indirect register addressing. Format: U12
	15:0	<b>Address subregister a0.4:uw.</b> This field, with only the lower 12 bits populated, can be used as an unsigned integer for register-indirect register addressing. Format: U12
1	31:16	<b>Address subregister a0.3:uw.</b> This field, with only the lower 12 bits populated, can be used as an unsigned integer for register-indirect register addressing. Format: U12
	15:0	<b>Address subregister a0.2:uw.</b> This field, with only the lower 12 bits populated, can be used as an unsigned integer for register-indirect register addressing. Format: U12
0	31:16	<b>Address subregister a0.1:uw.</b> This field can be used for register-indirect register addressing or serve as message descriptor for <i>send</i> instruction. When used for register-indirect register addressing, it is a 12-bit unsigned integer. For <i>send</i> instruction, it provides the higher 16 bits of <desc>. Format: U12 or U16.
	15:0	<b>Address subregister a0.0:uw.</b> This field can be used for register-indirect register addressing or serve as message descriptor for <i>send</i> instruction. When used for register-indirect register addressing, it is a 12-bit unsigned integer. For <i>send</i> instruction, it provides the lower 16 bits of <desc>. Format: U12 or U16.

When used as a source or destination operand, the address subregisters can be accessed individually or as a group. In the following example, the first instruction moves all 8 address subregisters to the first half of GRF register r1, the second instruction replicates a0.4:uw as an unsigned word to the second half of r1, the third instruction moves the first 4 words in r1 into the first 4 address subregisters, and the fourth instruction replicates r1.4 as an unsigned word to the last 4 address subregisters.

```

mov (8) r1.0<1>:uw a0.0<8;8,1>:uw // r1.n = a0.n for n = 0 to 7 in words
mov (8) r1.8<1>:uw a0.4<0;1,0>:uw // r1.m = a0.4 for m = 8 to 15 in words
mov (4) a0.0<1>:uw r1.0<4;4,1>:uw // a0.n = r1.n for n = 0 to 3 in words
mov (4) a0.4<1>:uw r1.4<0;1,0>:uw // a0.m = r1.4 for m = 4 to 7 in words

```

When used as the register-indirect addressing for GRF registers, the address subregisters can be accessed also individually or in group. When accessed in group, the address subregisters must be group-aligned. For example, when two address subregisters are used for register indirect addressing, they must



be aligned to even address subregisters. In the following example, the first instruction is legal. However, the second one is not. As ExecSize = 8 and the width of <src0> is 4, two address subregisters will be used as row indices, each pointing to 4 data elements spaced by HorzStride = 1 dword. For the first instruction, the two address subregisters are a0.2:uw and a0.3:uw. The two align to a dword group in the address register. However, the two address subregisters for the second instruction are a0.3:uw and a0.4:uw. They are not dword aligned in the address register and therefore violate the above mentioned alignment rule.

```
mov (8) r1.0<1>:d r[a0.2]<4,1>:d      // a0.2 and a0.3 is used for src1
mov (8) r1.0<1>:d r[a0.3]<4,1>:d      // ILLEGAL use of register indirect
```

**Implementation restriction:** ISA supports per channel indexing for a source operand. As there are only 8 sub-fields in the address register (to save hardware cost), the execution size of an instruction using per-channel indexing is limited to 8. Software may reload the address register and use compression control **SecHalf** to complete a 16-channel computation.

**Implementation restriction:** When used as the source operand <desc> for the send instruction, only the first dword subregister of a0 register is allowed (i.e. a0.0:ud, which can be viewed as the combination of a0.0:uw and a0.1:uw). In addition, it must be of UD type and in the following form <desc> = a0.0<0;1,0>:ud.

**Implementation restriction:** Elements a0.0 and a0.1 have 16 bits each, but the rest of elements (a0.2:uw through a0.7:uw) only have 12 bits populated each. 12-bit precision supports full indirect-addressing capability for the largest GRF register range. Software must observe the asymmetry of the implementation. When a0.0:uw and a0.1:uw are the source or destination, full 16-bit precision is preserved. However, when a0.2:uw to a0.7:uw are the destination, the higher 4 bits for each element will be dropped; when they are the source, hardware inserts zero to the higher 4 bits for each element.

**Performance Note:** There is only one scoreboard for the whole address register. When a write to some subregisters is in flight, hardware will stall any instruction writing to other subregisters. Software may use the destination dependency control {NoDDChk, NoDDClr} to improve performance in this case. Similarly, when a write to some subregisters is in flight, hardware will stall any instruction sourcing other subregisters until the write retires.

### 5.3.4.5 Accumulator Registers

ARF Register Type Encoding (RegNum[7:4]):	0010b
Number of Registers:	2
Default Value:	None
Normal Access:	RW
Elements:	8 or 16
Element Size:	Various
Element Type:	Various
Access Granularity:	Word
Write Mask Granularity:	N/A



SecHalf Control: Yes  
 Index-ability: No

There are two accumulator registers, *acc0* and *acc1*. They can be accessed either as explicit source and/or destination registers or as implied source and/or destination registers. To a programmer, each accumulator register may contain either 8 doublewords or 16 words of data elements. However, as shown in

Table 5-9, each data element may have higher precision with additional guard bits than that indicated by the numerical data type.

**Table 5-8. Register and Subregister Numbers for Accumulate Register**

RegNum[3:0]	SubRegNum[4:0]
0000 = <b>acc0</b> 0001 = <b>acc1</b> All other encodings are reserved.	Reserved: MBZ The accumulator subfields are individually addressable at word granularity. When an accumulator register is an explicit destination, it follows the rules for a destination register. If an accumulator is an explicit source operand, its register region must match with that of the destination register. The precision for Accumulator for floating point is the exact same as a regular GRF register.

The accumulators will be updated implicitly only if the **AccWrCtrl** is on for the current instruction. The accumulator Disable in control register cr0.0 allows software to over-write **AccWrCtrl** control for implicit accumulator update. The write enable in word granularity for the instruction will be used to update the Accumulator, the data in the disabled channels will not be updated.

When an accumulator register is used as an implicit source or destination operand, it is *acc0* by default. For a SIMD16 DW/float instruction, both *acc0* and *acc1* are used. If **SecHalf** is set, the implicit accumulator is then *acc1*.

It is illegal to specify different accumulator registers for source and destination operands in an instruction (e.g. “*add (8) acc1:f acc0:f*”). Result of such instruction is unpredictable.

For a SIMD16 DW/float instruction, if an accumulator register is used as an explicit source or destination operand, it must be *acc0*.

When an accumulator register is used as an explicit source operand, it must be the first source operand <src0>.

Whether the accumulator register is updated for a given instruction depends on several conditions: it can be an explicit destination operand, it can be an implicit destination if the **AccWrCtrl** is set in the instruction.

Swizzling is not allowed when accumulator is used as implicit source or explicit source of an instruction.

**Implementation Precision Restriction:** *As there are only 64 bits per channel in dword mode (D and UD), it is sufficient to store multiplication result of two dword operands as long as the post source modified sources are still within 32 bits. If any one source is type UD and is negated, the negated result*



becomes 33 bits. The dword multiplication results will be 65 bits, bigger than the storage capacity of accumulators. Consequently, the results are unpredictable.

**Implementation Precision Restriction:** As there are only 33 bits per channel in word mode (W and UW), it is sufficient to store multiplication result of two word operands with and without source modifier as the result is up to 33 bits. Integer is stored in accumulator in 2's compliment form with bit 32 as the sign bit. As there is no guard bit left, the accumulator can only be sourced once before running into risk of overflowing. When overflow occurs, only modular addition can generate correct result. But in this case, conditional flags may be incorrect. When saturation is used, the output is unpredictable. This is also true for other operations that may result in more than 33 bits of data. For example, adding UDW (FFFFFFFF) with DW (00000001) results in (1FFFFFFFFE). The sign bit is now at bit 34 and is lost when stored in the accumulator. When it is read out later from the accumulator, it becomes a negative number as bit 32 now becomes the sign bit.

**Table 5-9. Accumulator Channel Precision**

Data Type	# Channel	Bits / Channel	Description
F	8	32	When accumulator is used for float, it has the exact same precision as any GRF register
D (UD)	8	64	When the internal execution data type is doubleword integer, each accumulator register contains 8 channels of (extended) doubleword integer values. The data are always stored in accumulator in 2's compliment form with 64 bits total regardless of the source data type. This is sufficient to construct the 64-bit D or UD multiplication results using an instruction macro sequence consisting mul, mach and shr (or mov).
W (UW)	16	33	When the internal execution data type is word integer, each accumulator register contains 16 channels of (extended) word integer values. The data are always stored in accumulator in 2's compliment form with 33 bits total. This supports single instruction multiplication of two word source in W and/or UW format.
B (UB)	N/A	N/A	Not supported data type.



### 5.3.4.6 Flag Register [DevSNB]

ARF Register Type Encoding (RegNum[7:4]):	0011b
Number of Registers:	1
Default Value:	None
Normal Access:	RW
Elements:	2
Element Size:	16 bits
Element Type:	UW
Access Granularity:	Word
Write Mask Granularity:	Word
SecHalf Control:	Yes
Index-ability:	No

There is one flag register that consists of two 16-bit subregisters. Each flag subregister can be individually addressed. Each bit of a flag subregister corresponds to a data channel. (See Table 3-10 for details). Furthermore, each 16-bit subregister may be split to half with the **SecHalf** control in the instruction.

The two flag subregisters (*f0.0:uw* and *f0.1:uw*) can be used as the destination of the conditional modifier and can also be the source of the predication. As both predication flag source and conditional flag destination share the same instruction field, when both are enabled, they use the same flag subregister.

The values held in the individual bits of a flag subregister are the result of the most recent instruction which performed a condition-code evaluation with that flag register specified as the destination of the evaluation result. For example “*add.nz.f0.0 ...*” updates flag *f0.0* with the per-channel results of the not-zero condition. The flag subregister has per-bit write enables. When being updated as the secondary destination associated with conditional modifier, only the bits corresponding to the enabled channels in *EMask* are updated. The other bits in the flag subregister are unchanged.

The flag register is not updated for SEL instruction and conditional branch instructions with embedded compares.

The flag register as a whole or as two subregisters can also be an explicit source and/or destination operand.

**Table 5-10. Register and Subregister Numbers for Flag Register**

RegNum[3:0]	SubRegNum[4:0]
0000 = <i>f0:ud</i>	00000 = <i>f0.0:uw</i>
All other encodings are reserved.	00010 = <i>f0.1:uw</i>
	All other encodings are reserved.



**Table 5-11. Flag Register Fields**

DWord	Bit	Description
0	31:16	<b>Flag subregister f0.1:uw</b> Format: U16 Same as f0.0:uw. FormatDesc
	15:0	<b>Reserved</b>

### 5.3.4.7 State Registers

ARF Register Type Encoding (RegNum[7:4]):	0111b
Number of Registers:	1
Default Value:	Provided by the Dispatcher
Normal Access:	RW
Elements:	4
Element Size:	32 bits
Element Type:	UD
Access Granularity:	Byte
Write Mask Granularity:	N/A
SecHalf Control:	No
Index-ability:	No

Thread state registers are read-only registers. They can be accessed as a whole or individually using region-based addressing with byte granularity.

**Table 5-12. Register and Subregister Numbers for State Register**

RegNum[3:0]	SubRegNum[4:0]
0000 = <b>sr0</b> All other encodings are reserved.	Valid encoding range: 00000 – 00111 (in unit of byte) For example, 00000 = <b>sr0.0:ud</b> . It contains general thread states such as EUID/TID 00100 = <b>sr0.1:ud</b> . It contains the additional ID 01000 = <b>sr0.2:ud</b> . It contains the Dispatch Mask 01100 = <b>sr0.3:ud</b> . It contains the Vector Mask All other encodings are reserved.



**Table 5-13. State Register Fields**

DWord	Bits	Subfield Description
0 (sr0.0:ud)	31:28	Reserved: MBZ
	27:24	<b>FFID</b> (Fixed Function Identifier). Specifies which fixed function unit generates the current thread. This field is set at thread dispatch and is forwarded on the message bus for all out-going messages from this thread.
	23	<b>Priority Class.</b> This field, when set, indicates the thread belongs to the high priority class, which has higher scheduling priority over any thread with this field cleared. The priority field below determines the relative priority within the same priority class. This field is initialized by the thread dispatcher at thread dispatch time and stays unchanged throughout the life span of the thread.  This field is forwarded on the message bus to the message bus arbiter for all out-going messages. It serves as a priority hint for the target shared function. See shared function chapters whether and how a shared function uses this priority hint.  0 = Low priority class 1 = High priority class
	22:19	Reserved: MBZ
	18:16	<b>Priority.</b> This field is the relative aging priority of the thread. This field indicates the 'age' of this thread relative to other thread within the EU. No two threads within the same EU can have the same priority number (independent of the priority class value). Within the same priority class, an older thread (with a larger priority number) has higher schedule priority over a younger thread.  This field is set to zero at a thread's dispatch.  During a thread's run time, this field may or may not be incremented when a new thread is dispatched to the same EU. It is only incremented when another thread's priority number is incremented and reaches the same value. For example, if currently there is a thread with priority 0 on an EU, dispatching a new thread to this EU will cause the old thread's priority number being incremented to 1. However, if the active thread (assuming for simplicity there is only one) on an EU has a priority number 1 (or 2 or 3), dispatching a new thread to this EU will not change the old thread's priority number. As threads on an EU may terminate in arbitrary orders, the exact number for a thread depends on the dynamic execution of threads.
	15	Reserved: MBZ
	14	CID – The field specifies the Core Identifier for the the EU.
	13	SID – The field specifies the Slice Identifier for the the EU.
	12	HSID – The field specifies the HalfSlice Identifier for the EU.
	11:8	EUID[3:0]. The field specifies the Execution Unit Identifier. The MSB of this field specifies the RowID.
1 (sr0.1:ud)	7:3	Reserved: MBZ
	2:0	<b>TID</b> (The thread identifier). Specifies which thread slot the current thread is assigned to. This field is set at thread dispatch.
2 (sr0.2:ud)	31:24	<b>FFTID.</b> This field contains the thread ID assigned in fixed functions. There is no connection between FFTID and the TID assigned in the EUs.
	23:0	Reserved : MBZ
	31:0	<b>Dispatch Mask (DMask).</b> The 32-bit field specifies which channels are active at Dispatch time. This field is used by hardware to initialize the mask register.  Format: U32



DWord	Bits	Subfield Description
3 (sr0.3:ud)	31:0	<b>Vector Mask (VMask)</b> . The 32-bit field contains for each 4-bit group the OR of the corresponding 4-bit group in the dispatch mask. This field is used by hardware to initialize the mask register. Format: U32

### 5.3.4.8 Control Register

ARF Register Type Encoding (RegNum[7:4]):	1000b
Number of Registers:	1
Default Value:	Provided by the Dispatcher
Normal Access:	RW
Elements:	4
Element Size:	32 bits
Element Type:	UD
Access Granularity:	Dword
Write Mask Granularity:	Dword
SecHalf Control:	No
Index-ability:	No

The Control register is a read-write register. It contains four 32-bit subregisters that can be accessed individually.

Subregister *cr0.0:ud* contains normal operation control fields such as the floating point mode and the accumulator disable. It also contains the master exception status/control field that allows software to switch back to the application thread from the system routine.

Subregister *cr0.1:ud* contains the mask and status/control fields for all exceptions. The exception fields are arranged in significance-decreasing order from MSB to LSB. This allows the system routine to use *lzd* instruction to find the high priority exceptions and handles them first. As each exception is mapped to a single bit, other exception priority order may be implemented by software. System routine may choose to handle one exception at a time, by handle the exception detected by a *lzd* instruction and return to application thread. Or it may choose to handle all the concurrent exceptions, by looping through the exception fields until all outstanding exceptions are handled before returning back to the application thread.

Exception enable bits (bits 15:0 in *cr0.1:ud*) control whether an exception will cause hardware to jump to system routine or not. Exception status and control bits (bits 31:16 in *cr0.1:ud*) indicate which exceptions have occurred and are used for system routine to clear the exception. Even if a given exception is disabled, the corresponding exception status and control bit still reflects the status whether an exception event has occurred or not.

*cr0.2:ud* contains the **Application IP (AIP)** indicating the current thread IP when an exception occurs.





*cr0.3.ud* is reserved. Writing to this subregister is dropped; result of reading from this subregister is unpredictable.

Fields in Control registers also refer to a virtual register called **System IP (SIP)**. SIP is the virtual register holding the global System IP, which is the initial instruction pointer for the system routine. There is only one SIP for the whole system. It is virtual only from a thread's point of view, as it is not visible (i.e. not readable and not writeable) to the thread software executed on an EU. It can only be accessed indirectly by the hardware to response to exception events. Upon an exception, hardware performs some book keepings (e.g. saving the current IP into AIP) and then jumps to SIP. Upon finishing exception handling, the system routine may return back to the application by clearing the Master Exception Status and Control field in *cr0*, which will cause the hardware to load AIP to IP register. See STATE\_SIP command for how to set SIP.

**Table 5-14. Register and Subregister Numbers for Control Register**

RegNum[3:0]	SubRegNum[4:0]
0000 = <i>cr0</i> All other encodings are reserved.	00000 = <i>cr0.0.ud</i> . It contains general thread control fields 00100 = <i>cr0.1.ud</i> . It contains exception status and control 01000 = <i>cr0.2.ud</i> . It contains AIP.10100 ( <b>reserved</b> ) All other encodings are reserved.

**Table 5-15. Control Register Fields**

DWord	Bits	Subfield Description
0 ( <i>cr0.0:ud</i> )	31	<p><b>Master Exception State and Control.</b> This field is the master state and control for all exceptions. Reading a 0 indicates that the thread is in normal operation state and a 1 means the thread is in exception handle state. Upon an exception event, hardware sets this field to 1 and switch to SIP.</p> <p>Writing a 1 to this field has no effect. Writing a 0 to this field also has no effect if the previous value is 0. In both cases, the field keeps the previous value.</p> <p>If the previous value of this field is 1, software writing a 0 causes the thread to return to AIP. This transition is automatic – software does not have to move AIP to IP. The value of this field then stays as 0.</p> <p>This field is initialized to 0.</p> <p>0 = Indicate that the thread is in normal state 1 = Indicate that the thread is in exception state</p>
	30:16	Reserved: MBZ



DWord	Bits	Subfield Description
	15	<p><b>Breakpoint Suppress.</b> This field specifies whether breakpoint exception is suppressed or not. This field is normally set by software and cleared by hardware. If Master Exception Status and Control field is 1, this field is ignored by hardware.</p> <p>If Master Exception Status and Control field is 0 (i.e. not in system routine) and Breakpoint is enabled: If this field is set, breakpoint is temporally ignored (suppressed); Upon a breakpoint condition, the instruction is executed and this bit is automatically reset by hardware.</p> <p>This field is provided to prevent infinite loop of jumping to the system routine on a breakpoint condition. The system routine must set this bit (and also clear the corresponding status and control field) before returning to the application thread.</p> <p>This field has no effect when Breakpoint Enable bits is cleared.</p> <p>This field is initialized to 0.</p> <p>0 = Breakpoint exception is not suppressed 1 = Breakpoint exception is suppressed</p>
	14:9	Reserved : MBZ
	8	<b>Reserved</b>
	7	Reserved : MBZ
	6	Reserved: MBZ
	5:4	Reserved: MBZ
	3	<p><b>Vector Mask Enable (VME).</b> This field indicates DMask or Vmask should be used by EU for execution.</p> <p>This field is set by the Thread Dispatch..</p> <p>0 : Use Dispath Mask (DMASK) 1 : Use Vector Mask (VMASK)</p>
	2	<p><b>Single Program Flow (SPF).</b> Specifies whether the thread has a single program flow (SIMDn<sub>xm</sub> with m = 1) or multiple program flows (SIMDn<sub>xm</sub> with m &gt; 1). This field affects the operation of all branch instructions.</p> <p>In Single Program Flow mode, all execution channels branch and/or loop identically.</p> <p>This field is initialized by the Thread Dispatch.</p> <p>0: Multiple Program Flows 1: Single Program Flow</p> <p><b>Programming Restrictions:</b></p> <ul style="list-style-type: none"> <li>• <b>[DevSNB]</b> The <i>fork</i> instruction MUST be used to toggle SPF in the middle of a program to bring the PciPs to the ExIP. Program is not allowed to write directly into this bit.</li> <li>• Only H1/Q1/N1 are allowed in SPF mode.</li> <li>• <b>[Errata DevSNB]</b> IP modification using non-branch instructions are not allowed in SPF mode.</li> </ul> <p><b>Power Optimization:</b> If an entire shader doesn't do SIMD branching, driver can set the SPF to 1 to save power in HW.</p>



DWord	Bits	Subfield Description
	1	<p><b>Accumulator Disable.</b> This field controls the update of the accumulator by the instruction field <i>AccWrCtrl</i>. If this field is cleared, the accumulator is updated for all instructions with <i>AccWrCtrl</i> enabled. If set, the accumulator is disabled for all update operations, maintaining its value prior to being disabled. Setting this field has no effect if the accumulator is the explicit destination operand for an instruction.</p> <p>This field is initialized to 0.</p> <p>0 = Enable accumulator update 1 = Disable accumulator update</p> <p><i>Usage Notes:</i></p> <p><i>This control bit is primarily designed for the System Routine. That routine is not expected to use the accumulator, though it may need to use instructions which include implicit update of the accumulator. In order to use those instructions within the System Routine, but still preserving the accumulator contents upon return to the application kernel, the System Routine would either (a) save and restore the accumulator, or (b) prevent the accumulator from being unintentionally modified. This control bit has been added for the latter method.</i></p> <p><i>Software has the option to limit the setting of this control bit strictly within the System Routine. If, by convention, this bit is clear within application kernels, the System Routine can simply set the bit upon entry and clear it prior to returning control to the application kernel. This usage model would not necessarily require cr0.0 to be saved/restored in the System Routine. However, if by convention application kernels are permitted to set this bit, then the System Routine would be required to preserve the content of this bit.</i></p> <p>[DevSNB+] Reserved, MBZ.</p>
	0	<p><b>Floating Point Mode (FPMode).</b> This field specifies whether the current floating point operation mode is in IEEE standard mode or the alternative mode. It is used to control the floating operation of the Execution Unit. It is also forwarded on the message sideband for all out-going messages, for example, to control the floating point mode of the Extended Math unit or the Sampler unit. Software may modify this field to dynamically switch between the two floating point modes.</p> <p>This field is initialized by the Thread Dispatch.</p> <p>0 = IEEE floating point mode 1 = Alternative floating point mode</p>
1 (cr0.1:ud)	31	<b>Reserved</b>
	30	<p><b>External Halt Exception Status and Control.</b> This field indicates the External Halt exception. It is set by EU hardware upon receiving the broadcast External Halt signal. System routine should reset this field before returning to application routine in order to avoid infinite loop.</p> <p>This bit may be set or cleared by software.</p> <p>This field is initialized to 0.</p>
	29	<p><b>Software Exception Control.</b> This is the control field of software exception. Setting this field to 1 in application routine will cause an exception. Clearing this field in application routine has no effect. Upon entering system routine, the hardware maintains this field as one to signify software exception. System routine should reset this field before returning to application routine.</p> <p>This field may be set or cleared by software.</p> <p>This field is initialized to 0.</p>



DWord	Bits	Subfield Description
	28	<p><b>Illegal Opcode Exception Status.</b> This field, when set, indicates illegal opcode exception. The exception handle routine normally does not return back to the application thread upon an illegal opcode exception. Leaving this bit set, has no effect on hardware – if system software adversely returns to application routine leaving this field set, it doesn't cause any exception. This field should not be set by software or left set by system routine to avoid confusion.</p> <p>This field is initialized to 0.</p>
	27	<p><b>Stack Overflow Exception Status.</b> This field, when set, indicates stack overflow exception. The exception handle routine normally does not return back to the application thread upon this exception exception. Leaving this bit set, has no effect on hardware – if system software adversely returns to application routine leaving this field set, it doesn't cause any exception. This field should not be set by software or left set by system routine to avoid confusion.</p> <p>This field is initialized to 0.</p>
	23:16	Reserved: MBZ
	15	<p><b>Breakpoint Enable.</b> Specifies whether breakpoint exception is enabled or not.</p> <p>This field is initialized by the Thread Dispatcher.</p> <p>Format = ENABLED</p> <p>0 = Disabled</p> <p>1 = Enabled</p>
	14	<b>Reserved</b>
	13	<p><b>Software Exception Enable.</b> This field enables or disables the software exception. Enabling or disabling this field may allow host software to turn on/off certain features (such as profiling) without changing the kernel program.</p> <p>This field is initialized by the Thread Dispatcher.</p> <p>Format = ENABLED</p>
	12	<p><b>Illegal Opcode Exception Enable.</b> This field specifies whether illegal opcode exception is enabled or not. Illegal opcode exception includes illegal opcode and undefined opcode, caused by bad program or run time data corruption.</p> <p>This field is initialized by the Thread Dispatcher.</p> <p>Software should normally set it in the interface descriptor. Even though the mechanism is provided to disable illegal opcode exception, it should be used with extreme caution.</p> <p>Format = ENABLED</p>
	11	<p><b>Stack Overflow Exception Enable.</b> This field specifies whether stack overflow exception is enabled or not. The exception includes an overflow or an underflow in the stack space allocated for the thread.</p> <p>This field is initialized by the Thread Dispatcher.</p> <p>Software should normally set this in the interface descriptor.</p> <p>Format = ENABLED</p>
	10:0	Reserved: MBZ



DWord	Bits	Subfield Description										
2 (cr0.2:ud)	31:3	<p><b>Application IP (AIP).</b> This is the register storing the instruction pointer before an exception is handled. Upon an exception, hardware automatically saves the current IP into the AIP register, and then sets the <b>Master Exception State and Control</b> field to 1, which forces a switch to the System IP (SIP). AIP register may contain either the pointer to the instruction that causes the exception or the one after (such as mask stack overflow/underflow exceptions). This is shown in the following table, where IP is the instruction which generated the exception.</p> <table border="1"> <thead> <tr> <th>Exception Type</th> <th>AIP Value</th> </tr> </thead> <tbody> <tr> <td>Breakpoint</td> <td>IP</td> </tr> <tr> <td>External Halt</td> <td>n/a <sup>(1)</sup></td> </tr> <tr> <td>Software Exception</td> <td>IP + 1</td> </tr> <tr> <td>Illegal Opcode</td> <td>IP</td> </tr> </tbody> </table> <p>(1) External Halt exception is asynchronous and not associated with an instruction. When the system routine changes the Master Exception State and Control field from 1 to 0. Hardware restores IP from this register. This field is writable allowing returning IP to be altered after an exception handle.</p>	Exception Type	AIP Value	Breakpoint	IP	External Halt	n/a <sup>(1)</sup>	Software Exception	IP + 1	Illegal Opcode	IP
Exception Type	AIP Value											
Breakpoint	IP											
External Halt	n/a <sup>(1)</sup>											
Software Exception	IP + 1											
Illegal Opcode	IP											
	2:0	Reserved : MBZ										

**Implementation Restriction on Register Access:** When the control register is used as an explicit source and/or destination, hardware doesn't ensure execution pipeline coherency. Software must set the thread control field to **'switch'** for an instruction that uses control register as an explicit operand. This is important as the control register is an implicit source for most instructions. For example, fields like FPMODE and Accumulator Disable control the arithmetic and/or logic instructions. Therefore, if the instruction updating the control register doesn't set 'switch', subsequent instructions may have indeterminate results.

### 5.3.4.9 Notification Registers

ARF Register Type Encoding (RegNum[7:4]):	1001b
Number of Registers:	3
Default Value:	No
Normal Access:	RO
Elements:	3
Element Size:	32 bits
Element Type:	UD
Access Granularity:	Dword
Write Mask Granularity:	Dword
SecHalf Control:	No
Index-ability:	No



There are three notification registers (*n0.0:ud*, *n0.1:ud*, and *n0.2:ud*) used by the *wait* instruction. These registers are read-only and can be accessed in 32-bit granularity.

It should be noted that in the extreme case, it is possible to have more notifications to a thread than the maximal allowable of concurrent threads in the system. Therefore, the range of the thread-to-thread notification count in *n0*, is larger than the maximum number of threads computed by  $EUID * TID$ .

There is only one bit for the host-to-thread notification count in *n1*.

When directly accessed, this register is read-only. If the value is none zero, the only way to alter the value is to use the *wait* instruction to decrement the value until zero is reach. A *wait* instruction on a zero notification subregister will cause the thread to stall, waiting for a notification signal from outside targeting to the same subregister. See *wait* instruction for details.

*Implementation Restrictions: The notification registers are initialized to 0 after hardware/software reset. However, it is not reset at thread dispatch time.*

**Table 5-16. Register and Subregister Numbers for Notification Register**

RegNum[3:0]	SubRegNum[4:0]
0000 = <i>n0</i> All other encodings are reserved.	00000 = <i>n0.0:ud</i> 00100 = <i>n0.1:ud</i> 01000 = <i>n0.2:ud</i> All other encodings are reserved.

**Table 5-17. Fields of Notification Register *n0***

DWord	Bits	Subfield Description
0	31:16	Reserved: MBZ
	15:0	<b>Thread to Thread Notification Count.</b> This register is used by the WAIT instruction for thread-to-thread synchronization. The value read from this register specifies the outstanding notifications received from other threads. It can be changed indirectly by using the WAIT instruction. See WAIT instruction for details. Format: U16

**Table 5-18. Fields of Notification Register *n1***

DWord	Bits	Subfield Description
0	31:1	Reserved : MBZ
	0	<b>Reserved</b>


**Table 5-19. Fields of Notification Register *n2***

DWord	Bits	Subfield Description
0	31:16	Reserved: MBZ



DWord	Bits	Subfield Description
	15:0	<b>Thread to Thread Notification Count.</b> This register is used by the WAIT instruction for thread-to-thread synchronization. The value read from this register specifies the outstanding notifications received from other threads. It can be changed indirectly by using the WAIT instruction. See WAIT instruction for details. Format: U16

**Table 5-20. Format of the Notification Register**



B6898-01

### 5.3.4.10 IP Register

ARF Register Type Encoding (RegNum[7:4]):	1010b
Number of Registers:	1
Default Value:	Provided by the Dispatcher
Normal Access:	RW
Elements:	1
Element Size:	32 bits
Element Type:	UD
Access Granularity:	Dword
Write Mask Granularity:	Dword
SecHalf Control:	No
Index-ability:	No

The ip register can be accessed as a 32-bit quantity. It is a read-write register, containing the current instruction pointer, which is relative to the **Generate State Base Address**. Reading this register returns the instruction pointer of the current instruction. The 3 LSBs are always read as zero. Writing this register forces the program flow to jump to the new address. When it is written, the 3 LSBs are dropped by hardware.

[Errata **DevSNB{WA}**] - If IP is updated using a non-flow control instruction by a positive offset, the minimum positive offset must be 2.

[Errata **DevSNB{WA}**] - When SPF is ON, IP may not be updated by non-flow control instructions.



**Table 5-21. Register and Subregister Numbers for IP Register**

RegNum[3:0]	SubRegNum[4:0]
0000 = <i>ip</i> All other encodings are reserved.	00000 = <i>ip:ud</i> All other encodings are reserved.

**Table 5-22. IP Register Fields**

DWord	Bits	Subfield Description
0	31:3	<b>Ip.</b> Specifies the current instruction pointer. This pointer is relative to the <b>General State Base Address</b> .
	2:0	Reserved : MBZ

### 5.3.4.11 TDR Register

ARF Register Type Encoding (RegNum[7:4]):	1011b
Number of Registers:	8
Default Value:	No
Normal Access:	RO/CW
Elements:	8
Element Size:	16 bits
Element Type:	UW
Access Granularity:	Word
Write Mask Granularity:	Word
SecHalf Control:	No
Index-ability:	No

There are 8 thread dependency registers (tdr0.0:uw, tdr0.1:uw, tdr0.2:uw, tdr0.3:uw, tdr0.4:uw, tdr0.5:uw, tdr0.6:uw, and tdr0.7:uw) used by HW for the *sendc* instruction. These registers are read-only and can be accessed in 16-bit granularity.

When accessed explicitly, each thread dependency register has FFTID in the lower 8 bits, bits8 to bits14 are forced to zero by HW. Bit[15] is the valid bit, which indicate whether the current thread has a dependency on the dependency thread stored in this thread dependency register..

The thread dependency registers are read only, the valids can only be set with a thread dispatch, and it will reset by the broadcasting end of thread messages after a thread retired. The FFTID's can only be changed with a thread dispatch. Any write into any of the TDR register will clear the valid bit for the particular TDR if the write enable is true, the FFTID portion is strictly read only.





**Table 5-23. Register and Subregister Numbers for TDR Register**


RegNum[3:0]	SubRegNum[4:0]
1011 = <i>tdr0</i> All other encodings are reserved.	00000 = <i>tdr0.0:uw</i> 00001 = <i>tdr0.1:uw</i> 00010 = <i>tdr0.2:uw</i> 00011 = <i>tdr0.3:uw</i> 00100 = <i>tdr0.4:uw</i> 00101 = <i>tdr0.5:uw</i> 00110 = <i>tdr0.6:uw</i> 00111 = <i>tdr0.7:uw</i> All other encodings are reserved.

### 5.3.5 Immediate

Two forms of immediate are provided as a source operand: scalar and vector.


For a scalar immediate, it can be of any of the specified numerical data types from a word to a dword. Byte and unsigned byte are not supported as the smallest internal type of the execution pipeline is word. These two numerical types are reserved for future extensions.

The immediate field in an instruction has 32 bits as shown below. For a word or an unsigned word immediate data, software must replicate the same 16-bit immediate value to both the lower word and the high word of the 32-bit immediate field in an instruction.



The immediate form of vector allows a constant vector to be in-lined in the instruction stream. Both integer and float immediate vectors are supported.

An immediate integer vector is denoted by type *v* as *imm32:v*, where the 32-bit immediate field is partitioned into 8 4-bit subfields. Each 4-bit subfield contains a signed integer value in 2's complement form (halfbyte). Therefore each 4-bit subfield has a range of [-8, +7]. This is depicted in the following table.





An immediate float vector is denoted by type **vf** as *imm32:vf*, where the 32-bit immediate field is partitioned into 4 8-bit subfields. Each 8-bit subfield contains a signed float with 3-bit exponent and 4-bit fraction. Each 8-bit subfield provides signed floating point values with restricted range and precision. This is depicted in the following table.

31	24	23	16	15	8	7	0
immVF3			immVF1		immVF1		immVF0

*Restriction: When an immediate vector is used in an instruction, the destination must be 128-bit aligned with destination horizontal stride equivalent to a word for an immediate integer vector (**v**) and equivalent to a dword for an immediate float vector (**vf**).*

### 5.3.6 Region Parameters

Unlike conventional SIMD architectures where an N-bit wide SIMD instruction can only operate on N-bit aligned SIMD data registers, a region-based register addressing scheme is employed in the architecture. The region-based register addressing capability significantly improves the SIMD computation efficiency by providing per-instruction-based multiple data gathering from register file. This avoids instruction overhead to perform data pack, unpack, and shuffling, which has been observed on other SIMD architectures. One benefit of such capability is allowing various kinds of 3D Graphics API Shader compute models to run efficiently. Another benefit is allowing high throughput of media applications, which tend to operate on byte or word data elements.

This can be illustrated by the example shown in Figure 5-7 and Figure 5-8. As shown in Figure 5-7, a sequence of SIMD instruction is executed on a conventional load/store based superscalar machine with SIMD instruction extension. The data parallelism can be achieved by first level of loop unrolling. As shown, there is a second level of loop for the task. Before a given SIMD compute instruction, *Process (i)*, can proceed, there might be a load, a data rearrange and a data unpack (and conversion) instruction to load and prepare the input data. After the compute instruction is complete, it might also require pack, rearrange and store instructions, to format and save the same to memory. At the loop, other scalar computations such as loop count and address generation may be needed. For the same program, when the data can fit in the large GRF register file, the outer loop may be unrolled. Here one or a few block loads (using *send* instruction) may be sufficient to move the working set into GRF. Then the data shuffle can be combined with the processing operation with region-based addressing capability. Per operand float type and mixed data type operation may also allow DevSNB to combine data conditioning operations with computing operations. These techniques in the architecture help to achieve high compute efficiency and throughput for graphics and media applications.

Figure 5-7. Conventional SIMD Instruction Sequence

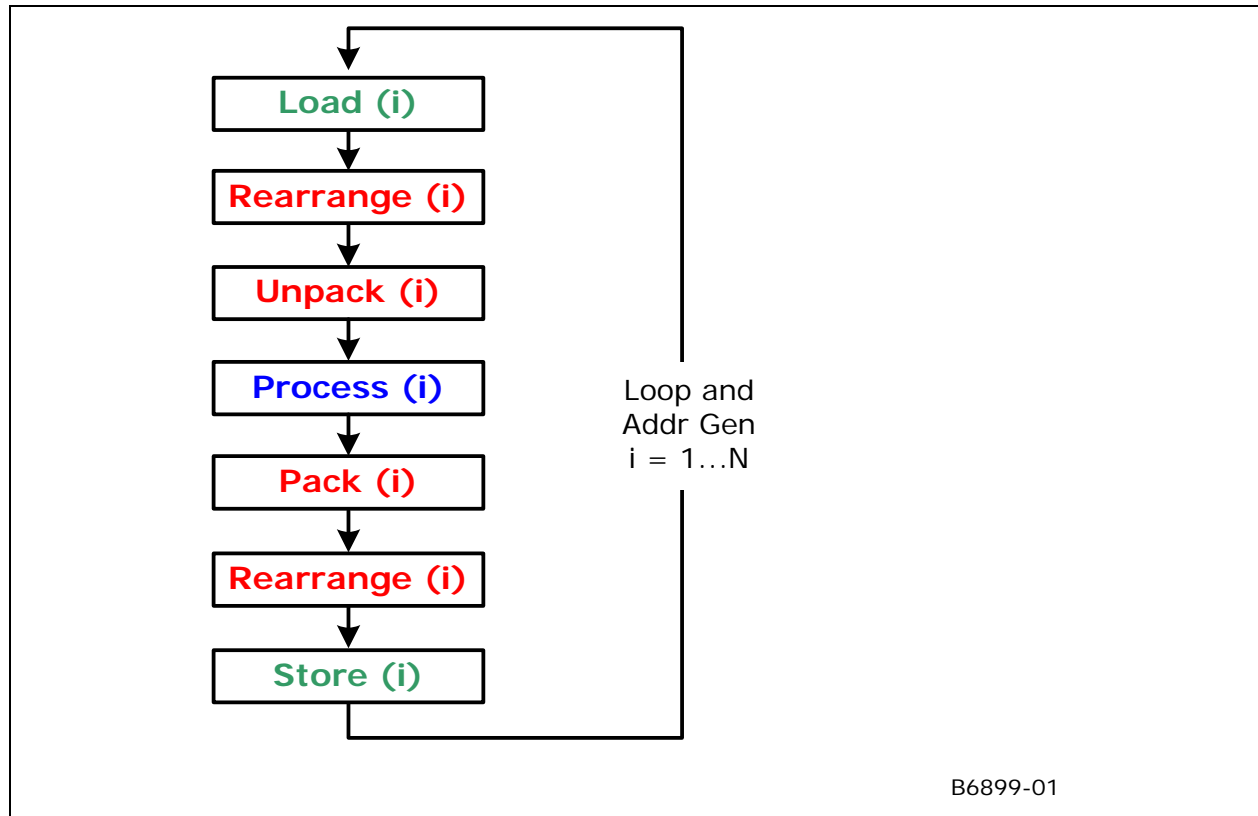
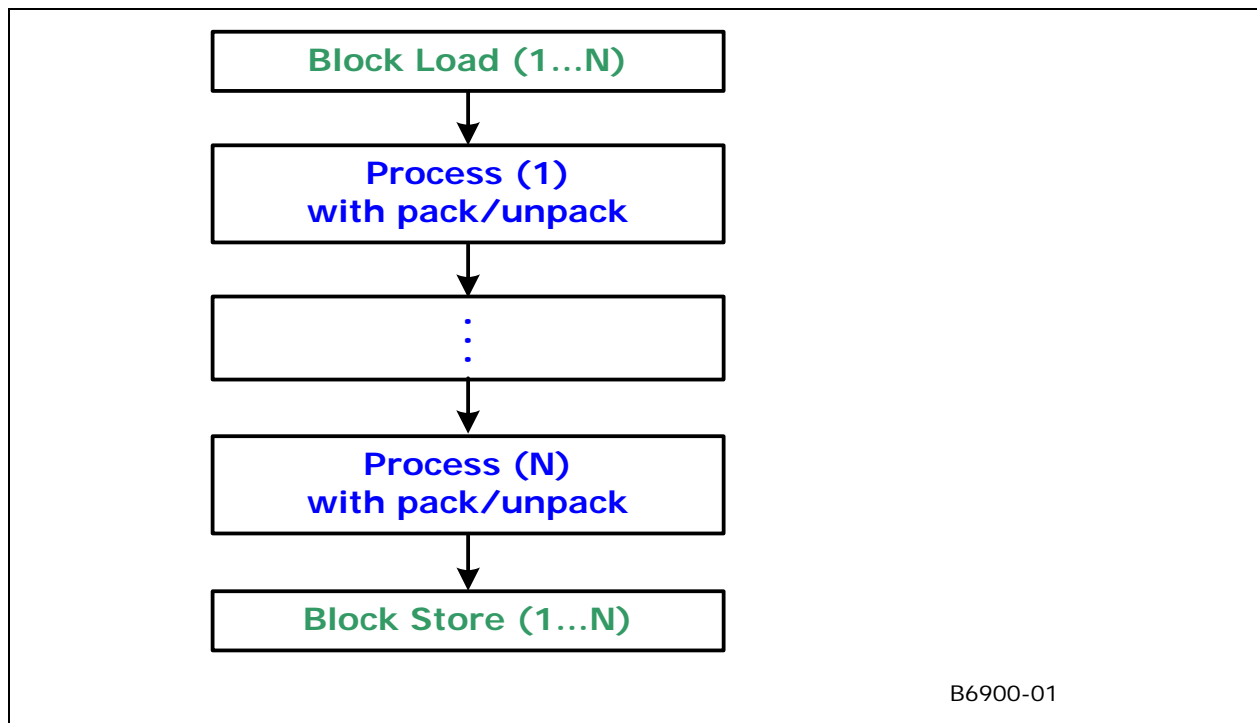




Figure 5-8. DevSNB SIMD Instruction Sequence for the Same Program



In an instruction, each operand defines a region in the register file. A region may contain multiple data elements. Each data element is assigned to an execution channel in the EU. The total number of data elements of a region is called the **size** of the region, or the size of the operand. The number of execution channels is called the **execution size (ExecSize)**, which is specified in the instruction word. ExecSize determines the size of region for source and destination operands in an instruction.

- For an instruction with two source operands, the sizes of the two source operands must be the same.
- The size of a destination operand generally matches the execution size, therefore equals to the number of source operand(s) in the same instruction.
  - Exception of this rule is present for the integer reduction instructions (such as sad2 and sada2) where the destination area is smaller than the source area.

Regions are **generalized 2-dimensional (2D)** arrays in row-major order. The first dimension is named the **horizontal** dimension (data elements within a row) and the second dimension is termed the **vertical** dimension (data elements in a column). Here, horizontal/vertical and row/column are just symbolic notations. When the GRF or MRF registers are viewed as a row-major 2D array of memory, such a notation normally matches well with the geometric locations of the data elements of an operand. However, as the register region is fully described by the parameters discussed below, the data elements of a register region may not form a regular rectangular shape. For example, Vertical Stride parameter is allowed to be smaller than Horizontal Stride, making the rows of a register region interleave with each other. It should also note that the meanings of horizontal/vertical here is different than that used for the flag control in Section 5.3.4.7.

Specifically, a region-based description of a source operand can take the following format

*RegFile RegNum.SubRegNum<VertStride;Width,HorzStride>:type*



Parameters are as the follows.

- Register Region Origin (*RegFile*, *RegNum* and *SubRegNum*): This set of parameters, including the register file, *RegFile*, the register number, *RegNum*, and the subregister number, *SubRegNum*, describes the register region origin, which is the location of the first data element of the operand. *RegNum* is in unit of 256-bit and *SubRegNum* is in unit of the data element size.
- Width (*Width*): *Width* specifies the number of data elements along the horizontal dimension, or the number of data elements of a row.
- Horizontal Stride (*HorzStride*): *HorzStride* specifies the step size between two adjacent data elements within a row. It is in unit of data element size, which is determined by the data element *Type*.
- Vertical Stride (*VertStride*): *VertStride* specifies the step size between two adjacent data elements along the vertical dimension (or the step size between two rows). It is again in unit of data element size, which is determined by the data element *Type*.
- Data Element Type (*Type*): *Type* specifies numerical data type (float, word, byte, etc.) of the data elements. All data elements within a region must have the same type.

Both GRF and MRF register files consist of a sequence of 256-bit physical registers. When viewing the register file (GRF for example) as a sequence of 256-bit aligned physical registers, *RegNum* field provides the physical register number, thus for the name. *SubRegNum* provides the sub-field addressing within a physical register. However, when viewing the register file as a byte addressable memory array, (*RegNum* and *SubRegNum*) is just a byte address within the register file with *SubRegNum* providing the lower 5 bits and *RegNum* providing the higher bits.

The execution size is specified for each instruction by the parameter *ExecSize*. The size of the vertical dimension is  $ExecSize/Width$ , based on the rule that the size of regions must equal to the execution size.

Figure 5-9 depicts the register region description. The example shows a register region of  $r4.1<16;8,2>:w$ , where the shaded fields denote the data elements in the region and the numbers in these fields are the execution channel assignments. The register region assumes that an *ExecSize* of 16 is set for the instruction. Each data element is a word (as noted by the type field “:w”). The origin of the region is at the second word of r4, denoted by *r4.1*. Each row of the region has 8 data elements (words) that are 2 data elements (words) apart. The distance between two rows is 16 words. Note that the region shown is for illustration purpose only. It does not represent a typical usage model nor a performance optimized mode.



Figure 5-9. An example of a register region ( $r4.1<16;8,2>:w$ ) with 16 elements


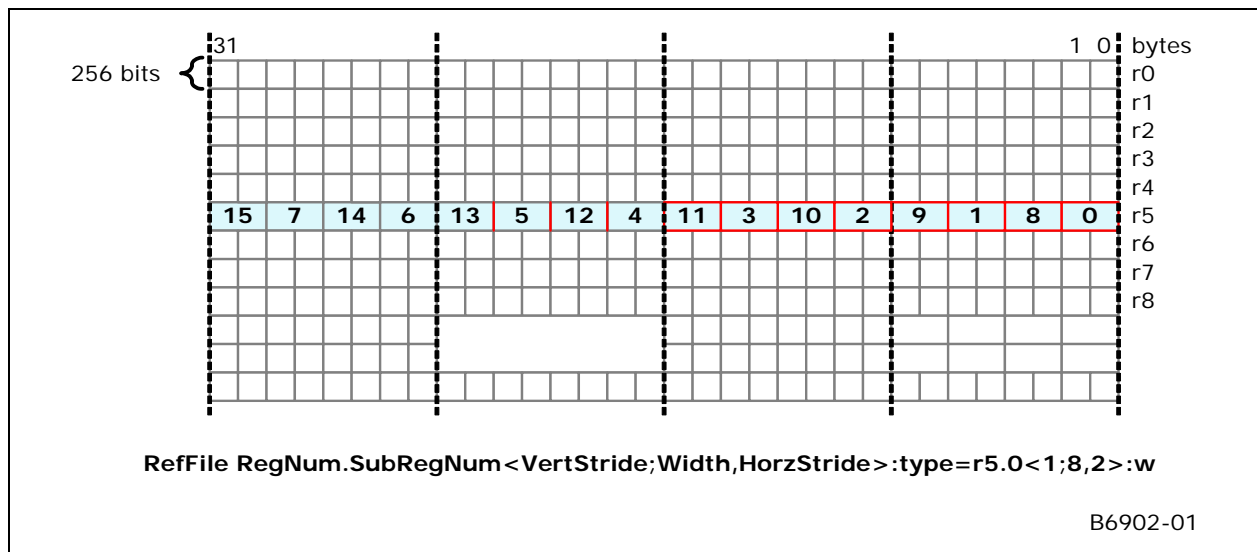


Figure 5-10 shows another example where the rows are interleaved. The region, having word data elements, starts at location  $r5.0:w$ . HorzStride, the distance within a row, is 2 words. So the second element (channel number 1) is at location  $5.2:w$ . And there are 8 elements per row. VertStride, the distance between two rows, is only 1 word, which is less than HorzStride. Therefore, the first element of the second row (channel number 8) is at  $r5.1:w$ , just next to channel number 0. It is clear from the picture that the two rows are interleaved.

By varying the region parameters, reader may construct other configurations. The next section provides more details on the region-based register addressing. However, there are restrictions imposed by hardware implementation, which can be found in the later sections of this chapter.



Figure 5-10. A 16-element register region with interleaved rows ( $r5.0<1;8,2>:w$ )



Without considering the source channel swizzle and destination register region description, the above row-major-order region description provides the data assignment to each execution channel. The following pseudo code computes the addresses of data elements assigned to execution channels for a special case when the destination register is aligned to 256-bit register boundary.

```
// Input:          Type: ub | b | uw | w | ud | d | f | v
//                RegNum: In unit of 256-bit register
//                SubRegNum: In unit of data element size
//                ExecSize, Width, VertStride, HorzStride: In unit of data elements
// Output:         Address[0:ExecSize-1] for execution channels

int ElementSize = (Type=="b"||Type=="ub") ? 1 : (Type=="w"||Type=="uw") ? 2 : 4;
int Height = ExecSize / Width;
int Channel = 0;
int RowBase = RegNum<<5 + SubRegNum * ElementSize;
for (int y=0; y<Height; y++) {
    int Offset = RowBase;
    for (int x=0; x<Width; x++) {
        Address [Channel++] = Offset;
    }
}
```



```
        Offset += HorzStride*ElementSize;
    }
    RowBase += VertStride * ElementSize;
}
```

As *HorzStride* and *VertStride* are specified independently (note that *VertStride* might be smaller than or equal to *HorzStride*), the region may take various shapes from a replicated scalar, a replicated vector, a vector of replicated scalars, a sliding window, to a non-overlapped 2D array.

A region-based description of a destination operand can take the following simplified format

```
RegFile RegNum.SubRegNum<HorzStride>.type
```

The destination operand is only allowed to have a 1 dimensional region. The Register Region Origin and Type are the same as for a source operand. The total number of elements is given by *ExecSize*. However, only *HorzStride* is required to describe the 1D region, not *VertStride* and *Width*.

As a source register region may across multiple physical GRF register, an instruction with such source operands may take more than two execution cycles to gather source data elements for execution. The destination register region is restricted to be within a physical GRF register. In other words, destination scatter writes over multiple physical registers are not supported.

## 5.3.7 Region Addressing Modes

There are two different register addressing modes: Direct register addressing and register-indirect register addressing. Depending on the register region description, the register-indirect register addressing mode can be further divided into three usages: 1x1 index region where only the origin of register region is provided by the address register, Vx1 index region where the offset of each row of the register region is provided by an address register, VxH index region where the offset of each data element is provided by an address register.

### 5.3.7.1 Direct Register Addressing

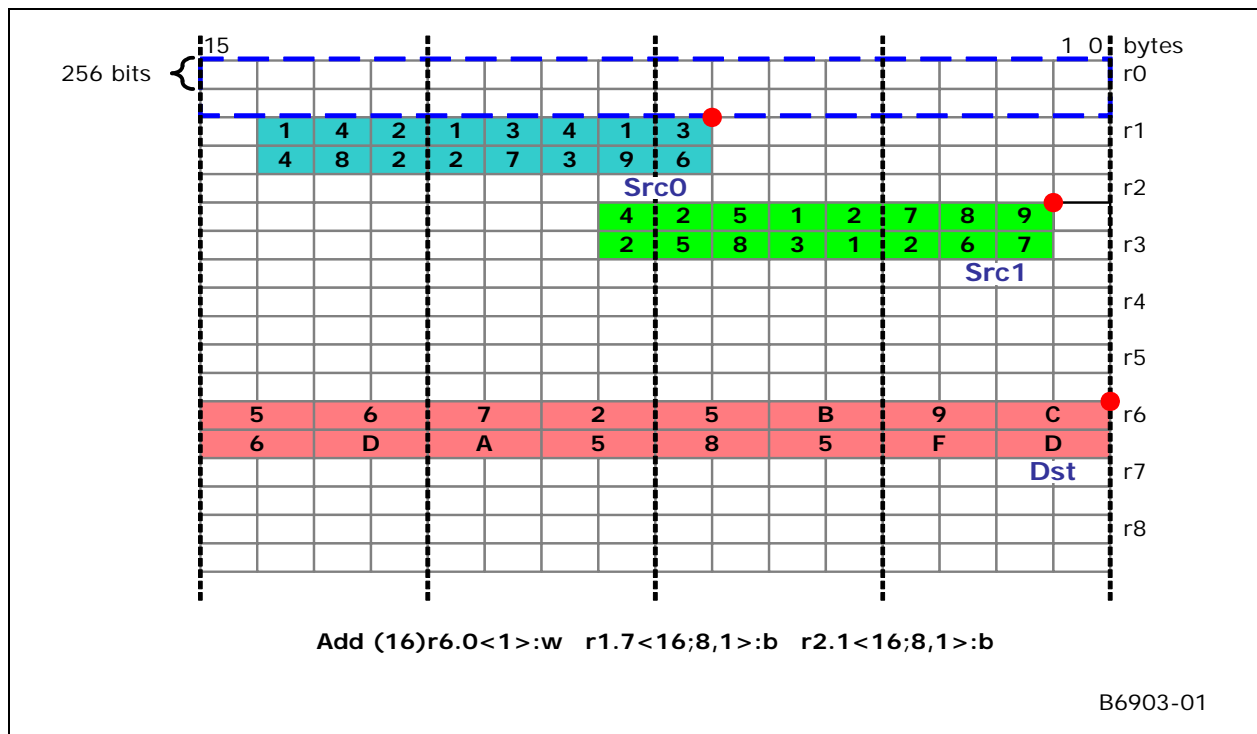
In this mode, all register region parameters are specified for an operand using fields in the instruction word.

Figure 5-11 and Figure 5-12 are two examples of direct register addressing.

For the example in Figure 5-11, all operands are 2D rectangular regions having the same size of 16 data elements. The two source operands, *Src0* and *Src1*, have 16 bytes. The destination operand, *Dst*, has 16 words. There are 8 elements in a row for *Src0* and *Src1*. The vertical stride of 16 bytes for *Src0* and *Src1* indicates that the first element and the 9'th element are 16 bytes apart in the register file. Note that *Src0* falls into the 256-bit physical GRF register starting at r1.0, but *Src1* crosses the 256-bit physical GRF register boundary between r2 and r3. The numbers in the shaded regions are the values of the data elements. Observing the upper right corners of the source/destination regions (first data element), we have  $C = 3+9$ .




Figure 5-11. A region description example in direct register addressing



For the example in Figure 5-12, the sizes of areas of *Src0* and *Src1* are the same, but *Src0* contains a vector of replicated scalars. With *HorzStride* = 0 and *Width* = 8, the first row of 8 elements in *Src0* is a replication of the byte at r1.14. Comparing *ExecSize* of 16 to *Width* of 8 indicates that there is a second row of 8 elements in *Src0*. With *VertStride* = 16, the second row in *Src0* is a replication of the byte at r1.20 (20 = 14+16). Effectively, the 16 data elements of *Src0* are {1,1,1,1,1,1,1,1, 4,4,4,4,4,4,4,4}.



Figure 5-12. A region description example in direct register addressing with <src0> as a vector of replicated scalars




### 5.3.7.2 Register-indirect Register Addressing with a 1x1 Index Region

In the register-indirect register addressing mode with 1x1 index region, the region origin is provided by the content of the address register, the rest of region parameters are provided by the fields in the instruction word.

Figure 5-13 depicts an example for this addressing mode. For example, the present of full region description <16;8,1> for Src0 indicates that only the origin of the region is provided by the address register a0.0.

**Figure 5-13. An example illustrating register-indirect register addressing mode with a 1x1 index region**




### 5.3.7.3 Register-indirect Register Addressing with a Vx1 Index Region

In the register-indirect register addressing mode with Vx1 index region, horizontal dimension is described by the fields in the instruction word and the vertical dimension is described by an address register region. Specifically, the origin of each row of the data region is provided by the contents of an address register region. The rows are described by the width and the horizontal stride. The first address register is provided, the following contiguous address registers are for the following rows. The total number of address registers used is inferred by parameters *ExecSize* and *Width*.

An example is provided in Figure 5-14. The assembly syntax notion of a register region without vertical stride, <4,1>, corresponding to the special encoding of vertical stride of 0xF in the instruction word, indicates the VxH or Vx1 mode of indirect register addressing. In this case, the origin for each row of Src0 is provided by the address register. As *ExecSize/Width* = 2, there are two address registers a0.0 and a0.1, each pointing to a row of 4 data elements.



Figure 5-14. An example illustrating register-indirect-register addressing mode with a Vx1 index region (Src0)




### 5.3.7.4 Register-indirect Register Addressing with a VxH Index Region

In the register-indirect register addressing mode with VxH index region, the position of each data element is provided by the contexts in an address register region. This mode has the identical syntax as the Vx1 index region mode, and in fact, can be viewed as a special case of the Vx1 mode. When *Width* of the region is 1, the number of address registers used equals *ExecSize*.

An example is provided in Figure 5-15. The absent of vertical stride in the region description <1,0> with width = 1 indicates that the origin for each row of 1 data element of Src0 is provided by the address register. As  $ExecSize/Width = 8$ , there are 8 address registers from a0.0 to a0.7, each pointing to a single data elements.

Figure 5-15. An example illustrating register-indirect register addressing mode with a VxH index region (Src0).



### 5.3.8 Access Modes

There are two basic register access modes controlled by a single bit instruction subfield called Access Mode.

- 16-byte Aligned Access Mode (**align16**): In this mode, the origins of all operands (sources and destination), whether it is by direct addressing or register-indirect addressing, are 16-byte aligned. For example a row in the region description starts at 16-byte aligned and the width the row must be 4 and the 4 data elements within a row must span 16-bytes. In this access mode (and with other restrictions put forward later), full-channel swizzle for both source operands and per-channel mask for destination operand are supported on a 4-component basis. In other words, the control and setting of full source swizzle and destination mask are repeated for every 4 components up to total of *ExecSize* channels.
  - The **align16** access mode can be used for AOS operations. See examples provided in the Primary Usage Model section for SIMD4x2 and SIMD4x1 modes of operation to support 3D API Vertex Shader and Geometric Shader execution.



- 1-byte Aligned Access Mode (**align1**): In this mode, the origins of all operands may be aligned to their data type and could be 1-byte if the operand is of byte type. In this access mode, full region register descriptions are supported, however, source swizzle or destination mask are not supported.
  - The **align1** access mode can be used for SOA operations. See examples provided in the Primary Usage Model section for SIMD8 and SIMD16 modes of operation to support 3D API Pixel Shader. Many media applications also operate well in **align1** access mode.

### 5.3.9 Execution Data Type

The architecture supports instructions with mixed data types. The internal hardware computation is performed using the execution data type. When an instruction has only one source operand or has two source operands of the same data type, the execution data type is the same as that of the source. When an instruction has two source operands of different types, an execution data type is determined and one of the source operands will be converted to the execution type before the computation is performed. The execution type is independent of the destination data type. When the destination data type is different from the execution data type, a type conversion is performed on the intermediate compute results before the results are written into the destination register. Such a destination type conversion also apply to accumulator registers, implicitly or explicitly.

Determination of the execution data type for two sources of different data types obeys the following rules

- Instruction with mixed float and integer type sources is not allowed.
- Else if any source is a dword, the execution data type is signed dword integer (D)
- Else execution data type is signed word integer (W)

Note that when the execution data type is an integer, it is always a signed integer. This doesn't affect the functional correctness of the instruction as extra precisions are carried within the hardware, including the accumulator. See Instruction Reference Chapter for detailed description for each instruction.

#### 5.3.10 Register Region Restrictions

The following register region rules apply to the implementation. Rules and restrictions for compressed instructions can be found in the Instruction Compression section.

1. *ExecSize* must be equal to or less than the maximum execution size supported for the operand type. As shown in Table 5-24, the maximum execution size is determined by the largest operand type of the sources and destination of the instruction.
2. The mapping of data elements within the region of a source operand is in row-major order and is determined by the region description of the source operand, plus *ExecSize* and destination region description.
3. *ExecSize* must be equal to or greater than *Width*.
4. If  $ExecSize = Width$  and  $HorzStride \neq 0$ , *VertStride* must be set to  $Width * HorzStride$ .
5. If  $ExecSize = Width$  but  $HorzStride = 0$ , there is no restriction on *VertStride*.



6. If *Width* = 1, *HorzStride* must be 0 regardless of the values of *ExecSize* and *VertStride*.
7. If *ExecSize* = *Width* = 1, both *VertStride* and *HorzStride* must be set to zero.
8. If *VertStride* = *HorzStride* = 0, *Width* must be 1 regardless of the value of *ExecSize*.
9. Destination region cannot cross the 256-bit register boundary.
  - 9.1. Exception to this rule is for a SIMD16 DW/Float instruction where the destination region covers exactly **two adjacent** 256-bit physical registers.
10. Destination region alignment rule.
  - 10.1. With the exception on 'raw move' described in rule #10.3 and the exception on byte destination in rule #10.5, all destination data elements must be aligned to the size for the execution data type of the instruction. For example, if one of the source operands is in dword mode (a float, a signed or unsigned dword integer), the execution data type will be either float or signed dword integer. Therefore, the destination data elements must be dword aligned. This rule has the following two implications:
    - 10.1.1. The destination sub-register must be aligned to the size of the execution data type.
    - 10.1.2. If *ExecSize* is greater than 1, *dst.HorzStride*\*sizeof(*dst.Type*) must be equal to or greater than the size of the execution data type.
  - 10.2. If *ExecSize* is 1, *dst.HorzStride* must not be 0. Note that this is relaxed from rule 10.1.2. Also note that this rule for destination horizontal stride is different from that for source as stated in rule #7.
  - 10.3. When destination type is byte (UB or B), only a 'raw move' using *mov* instruction supports packed byte destination register region: *dst.HorzStride* = 1 and *dst.type* = (UB or B). This packed byte destination region is not allowed for any other instructions, including a 'raw move' using *sel* instruction. This is because *sel* instruction is based on word or dword wide execution channels.
  - 10.4. When an instruction has a source region that spans two physical registers and destination register contained in one register, one of the followings must be true:
    - 10.4.1. Destination region is entirely contained in the lower oword of a physical register,
    - 10.4.2. Destination region is entirely contained in the upper oword of a physical register, or
    - 10.4.3. Destination elements are evenly split between the two owords of a physical register.
  - 10.5. When an instruction has a source region that spans two physical registers, the destination spans two physical registers, and the destination elements are evenly split between the two physical registers. Then each destination register must be entirely derived from one source register.
  - 10.6. Relaxed alignment rule for byte destination. When destination type is byte (UB or B), destination data elements can be either aligned to the lowest byte or the second lowest byte of the execution channel. For example, if one of the source operands is in word mode (a signed or unsigned word integer), the execution data type will be signed word integer. In this case, the destination data bytes can be either all in the even byte locations or all in the odd byte locations. This rule has the following two implications:



10.6.1. The destination sub-register must be either aligned to the size of the execution data type or one byte higher off the execution data type.

10.6.2. If *ExecSize* is greater than 1, *dst.HorzStride\* sizeof(dst.Type)* must be equal to or greater than the size of the execution data type. This is the same as that in #10.1.2.

11. In Align1 access mode, a source region must be within **two adjacent** 256-bit physical registers.

12. Rules on register-indirect register access:

12.1. An indexed source1 can only have a 1x1 indexed register region – only single index mode is allowed for a source 1.

12.2. An indexed destination can only have a 1x1 indexed register region – only single index mode is allowed for a destination operand.

12.3. Data elements referenced by a single index within a source region cannot cross 256-bit physical register boundary. This applies to register region with a single index or with multiple indices.

12.3.1. A register region with multiple indices may access multiple physical registers as long as data elements associated with each index follow the above-mentioned rule. For example instruction “mov (16) r0.0:uw r[a0.0]<2,2>:uw” is allowed. This is a source gathering instruction whereas the source operand may potentially touch 8 different physical GRF registers.

**Table 5-24. Execution size in device hardware**

Device	Native Instructions		
Max Operand Size	DWORD	WORD	BYTE
[DevSNB]	16	32	32

**Table 5-25. Indirect source addressing support available in device hardware**

Device	Indirect Source 0	Indirect Source 1
[DevSNB]	Yes	Yes

### 5.3.10.1 Examples

Some examples are provided here to illustrate the cases when the register region restrictions are violated. It is provided as informative material to help understanding these restrictions.

*Example 1:* The following instructions are illegal as they violate rule #10.1, as the destination is not aligned to the execution data type.

```
mov (1) r0.1<1>:b r2.0:w // dst.SubReg must be even
```

```
mov (2) r0.0<1>:b r2.0:w // dst.HorzStride must be >= 2
```





```
mov (2) r0.0<2>:b r2.0:d // dst.HorzStride must be >= 4
mov (2) r0.0<2>:b r2.0:f // dst.HorzStride must be >= 4
mov (1) r0.2<1>:b r2.0:d // dst.SubReg must be dword aligned
```

*Example 2:* This instruction is illegal as it violates rule #10.1.2, as when *ExecSize* = 1, *dst.HorzStride* cannot be zero.

```
mov(1) r0.0<0>:b r0.0:d
```

*Example 3:* This instruction is illegal as it violates rule #Error! Reference source not found., as the source contains one row of 2 elements that spans physical register r2 and r3.

```
mov (2) r1.0:d r2.7<2;2;1>:d
```

### 5.3.10.2 Different Raw Moves

**Definition of Raw Move:** Raw move is an operation that moves data elements from source to destination without altering the bit fields of the data elements. It must use one of the move instructions such as *mov*, *sel*, *movi*. Arithmetic instruction that results in unaltered bit fields of the data elements are not treated as raw move. A raw move may subject to the execution channel enables by using prediction or being present in multi-channel branch code segment. Type conversion by definition cannot be used in a raw move. Therefore, source and destination operands must be of the identical data type. For example, if both source and destination are float, for an arithmetic instruction, denorm will be flushed to zero. However, for a raw move, denorm will be preserved.

**Definition of Byte Raw Move:** As the minimal execution channel type is word, when the destination stride is greater than one byte, each data element of the source can be mapped to one execution channel. This is referred to as Byte Raw Move. Byte Raw Move allows the destination to be byte aligned, in other words, allowing the destination to not align to execution channels. Byte Raw move subjects to execution channel enables.

**Definition of Packed-Byte Raw Move:** As the minimal execution channel type is word, when the destination stride is equal to one byte, two data elements of the source are mapped to one execution channel. This is referred to as Packed-Byte Raw Move. Packed-Byte Raw Move allows the destination to be byte aligned, in other words, allowing the destination to not align to execution channels. However, as the data elements are not mapped to execution channels, undefined results may occur if Packed-Byte Raw Move is mixed with execution channel enables. So for Packed-Byte Raw Move, *WECtrl* should be used when there are un-enabled channels within the execution size of the instruction.

## 5.3.11 Destination Operand Description

### 5.3.11.1 Destination Region Parameters

Based on the above restrictions, a subset of register region parameters are sufficient to describe the destination operand:



- Destination Register Origin
  - Destination Register Number and Destination Subregister Number for direct register addressing mode
  - A Scalar Destination Register Index for register-indirect-register addressing mode
- Destination Register 'Region' – Note that destination register region does not have full region description parameters
  - Destination Horizontal Stride

## 5.4 SIMD Execution Control

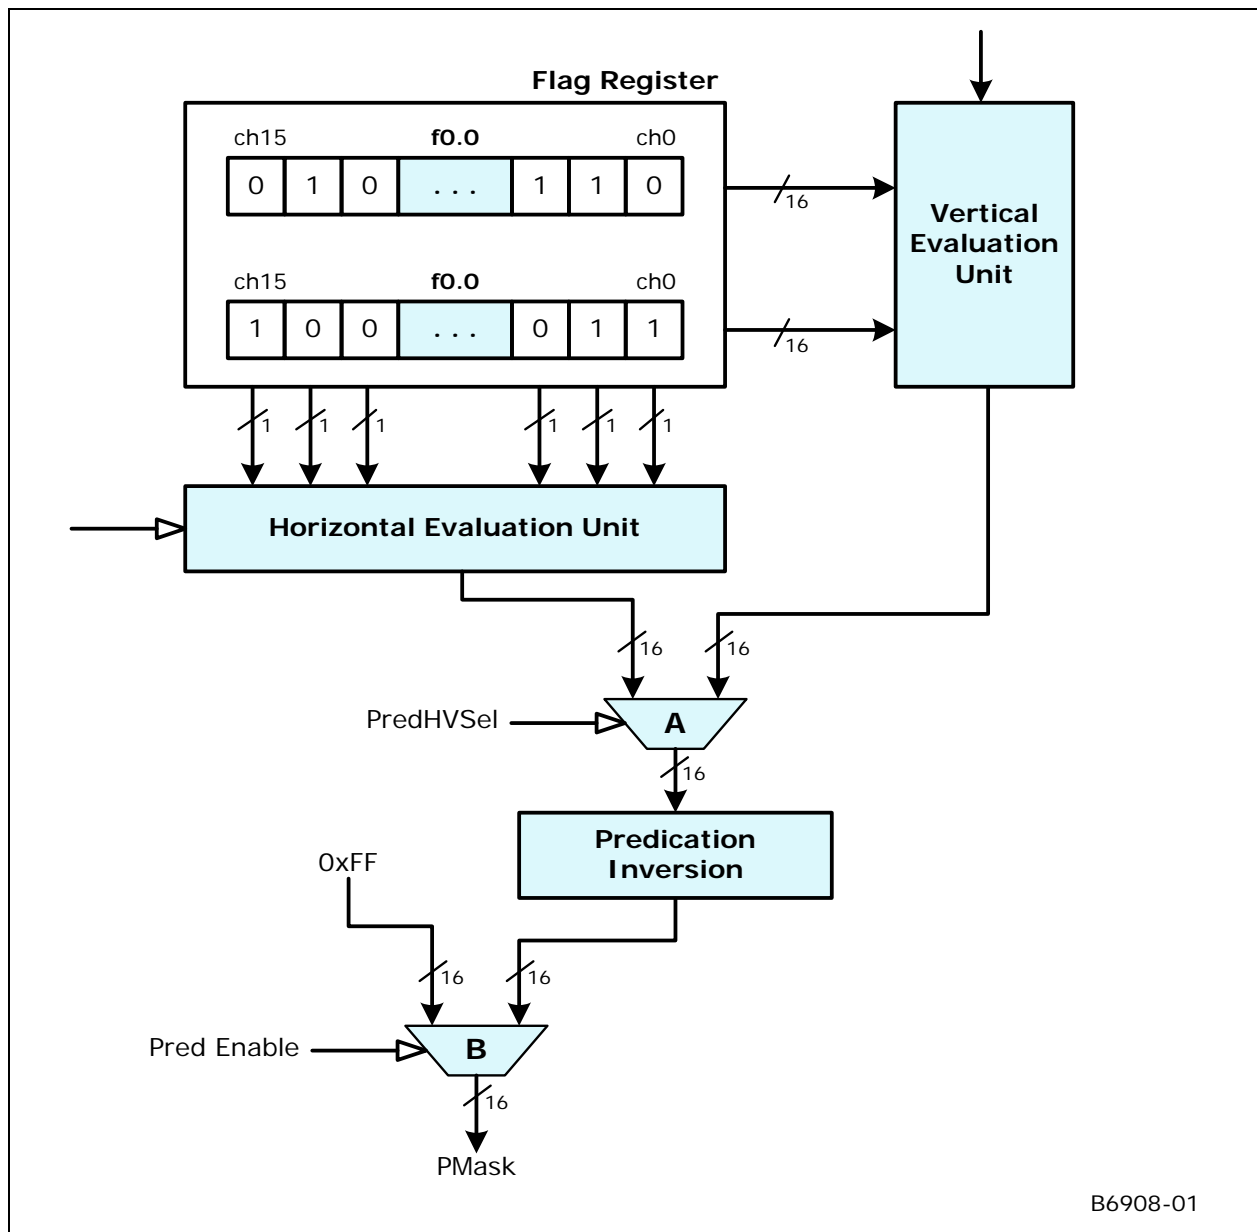
### 5.4.1 Predication

Predication is the conditional SIMD channel selection for execution on a per instruction basis. It is an efficient way of dynamic SIMD channel enabling without paying branch instruction overhead. When predication is enabled for an instruction, a Predicate Mask (PMask), which contains 16-bit channel enables, is generated internally in EU. Note that PMask is not a software visible register. It is provided here to explain how SIMD execution control works. PMask generation is based on the Predication Control (*PredCtrl*) field, Predication Inversion (*PredInv*) field and the flag source register in the instruction word. See Instruction Summary chapter for definition of these fields.

Figure 5-16 shows the block diagram of the hardware logic to generate PMask. PMask is generated based on combinatory logic operation of the bits in the flag register. Instruction field *PredCtrl* controls the horizontal evaluation unit and vertical evaluation unit. MUX A in the figure selects whether horizontally-evaluated results or vertically-evaluated results are sent to the Predication Inversion unit. The *PredInv* field controls the Prediction Inversion unit. Either one 16-bit flag subregister or the whole flag register may be selected to generate the PMask depending on the predication control modes. MUX B indicates that predication can be enabled and disabled. Predication can be grouped into the following three categories. Predication functionality also depends on the Access Mode of the instruction.

- No predication: Of course, predication can be disabled. This is the most commonly used case.
- Predication with horizontal combination: the predicate mask is generated based on combinatory logic operation of bits within a selected flag subregister.
- Predication with vertical combination: the predicate mask is generated based on combinatory logic operation of bits across flag multiple subregisters.

Figure 5-16. Generation of predication mask



### 5.4.2 No Predication

When PredCtrl field of a given instruction is set to 0 (“no predication”), it indicates that no predication is applied to this instruction. Effectively, the resulting PMask is all 1’s. This is shown by the 2:1 multiplexer B controlled by the Pred Enable signal in Figure 5-16. Where predication is not enabled for an instruction, multiplex B is selected to output 0xFF to PMask.



### 5.4.3 Predication with Horizontal Combination

Predication with horizontal combination inputs the 16 bits of a single flag subregister (f0.0:uw or f0.1:uw) and passes them through combinatory logic of the Horizontal Evaluation unit to create PMask.

The simplest combination is 'no combination' – the same 16 bits from selected flag subregister are output to MUX A. In this case, a bit in the selected flag subregister controls the conditional execution of the corresponding execution channel. Let the selected flag subregister be denoted as f0.#, the following pseudo code describes the predicate mask generation for predication with sequential flag channel mapping.

```
If (PredCtrl == "Sequential flag channel mapping") {  
    For (ch=0; ch<16; ch++)  
        PMask[ch] = (PredInv == TRUE) ? ~f0.#[ch] : f0.#[ch];  
}
```

More complex horizontal evaluation is based on channel grouping. A group of adjacent channels (bits from flag subregister) are evaluated together and a single bit is replicated to the group. The size of groups is in power of 2. The supported combination depends on the Access Mode of an instruction.

In **Align16** access mode, horizontal combination is based on 4-channel groups.

- Channel replication: PredCtrl of '.x', '.y', '.z' and '.w' select a single channel from each 4-channel group and replicate it as the output for the group. For example, PredCtrl = '.x' means that channel 0 in each group is replicated.
- OR combination: PredCtrl of '.any4h' means that if **any** of the channel in a group is enabled, outputs for the 4 channels in the group are all enabled.
- AND combination: PredCtrl of '.all4h' means that only when **all** of the channels in a group are enabled, the output for the group is enabled.

These combinations in **Align16** mode can be described by the following pseudo-code.

```
If (Access Mode == Align16) {  
    For (ch = 0; ch < 16; ch += 4)  
        Switch (PredCtrl) {  
            Case '.x':      bTmp = f0.#[ch]; break;  
            Case '.y':      bTmp = f0.#[ch+1]; break;  
            Case '.z':      bTmp = f0.#[ch+2]; break;  
            Case '.w':      bTmp = f0.#[ch+3]; break;  
            Case '.any4h':  bTmp = f0.#[ch] | f0.#[ch+1] | f0.#[ch+2] | f0.#[ch+3]; break;  
            Case '.all4h':  bTmp = f0.#[ch] & f0.#[ch+1] & f0.#[ch+2] & f0.#[ch+3]; break;  
        }  
}
```



```
bTmp = (PredInv == TRUE) ? ~bTmp : bTmp;
PMask[ch] = PMask[ch+1] = PMask[ch+2] = PMask[ch+3] = bTmp;
}
}
```

In **Align1** access mode, horizontal combination is based on AND combination ‘.any#h’ and OR combination ‘.all#h’ on channel groups with various sizes, where # is the number of channels in a group ranging from 2 to 16. This is described by the following pseudo-code.

```
If (Access Mode == Align1) {
    Switch (PredCtrl) {
        Case ‘.any2h’: groupSize = 2; <op> = ‘|’; break;
        Case ‘.all2h’: groupSize = 2; <op> = ‘&’; break;
        Case ‘.any4h’: groupSize = 4; <op> = ‘|’; break;
        Case ‘.all4h’: groupSize = 4; <op> = ‘&’; break;
        Case ‘.any8h’: groupSize = 8; <op> = ‘|’; break;
        Case ‘.all8h’: groupSize = 8; <op> = ‘&’; break;
        Case ‘.any16h’: groupSize = 16; <op> = ‘|’; break;
        Case ‘.all16h’: groupSize = 16; <op> = ‘&’; break;
    }
    For (ch = 0; ch < 16; ch += groupSize) {
        For (inc = 0, bTmp = FALSE; inc < groupSize; inc ++)
            bTmp = bTmp <op> f0.#[ch+inc];
        For (inc = 0; inc < groupSize; inc ++)
            PMask[ch+inc] = bTmp;
    }
}
```



## 5.4.4 Predication with Vertical Combination

Predication with vertical combination uses both flag subregister as inputs. The AND or OR combination is across the subregisters on a channel by channel basis. This is shown by the following pseudo-code.

```

If (Access Mode == Align1) {
    For (ch = 0; ch < 16; ch ++ ) {
        If (PredCtrl == 'any2v')
            PMask[ch] = f0.0[ch] | f0.1[ch]
        Else If (PredCtrl == 'any2h')
            PMask[ch] = f0.0[ch] & f0.1[ch]
    }
}

```

## 5.5 Instruction Compaction

### 5.5.1 Motivation and Expected Usage

Instruction Compaction is used to reduce the memory footprint of the shaders. It relies on predefined tables to compact instructions if they has patterns matching the entries inside the compaction table.

Because of the limited size of the compaction table in HW, not all instructions will be compacted by Jitter, HW receives both compacted and noncompactd instructions during execution.

Jitter compacts the instruction using the same compaction table inside HW, Jitter need to calculate any branch/jump offset after instruction compaction is done.

#### 5.5.1.1 Instruction Compaction [DevSNB]

The compaction table is listed below.

Control Bits		
0	0000000000000000	Align1  ,(1),
1	0100000000000000	Align1  ,(16),
2	0011000000000000	Align1  ,(8),
3	0000000010000000	Align1   xyzw, (1),
4	0001000000000000	Align1  ,(2),
5	0000100010000000	Align1 - xyzw, (1),
6	0000000010000010	Align1 NoMask   xyzw, (1),
7	0000000000000010	Align1 NoMask  ,(1),



Control Bits		
8	0100000010000000	Align1   xyzw ,(16),
9	0101000000000000	Align1  ,(32),
10	1011000000000000	Align1  ,(8), sat
11	0010000000000000	Align1  ,(4),
12	1101000000000000	Align1  ,(32), sat
13	1100000000000000	Align1  ,(16), sat
14	0100100010000000	Align1 - xyzw ,(16),
15	0100000000000100	Align1 NoDDChk  ,(16),
16	0100000000000100	Align1 NoDDClr  ,(16),
17	0000000000000100	Align1 NoDDChk  ,(1),
18	0000000000000100	Align1 NoDDClr  ,(1),
19	0011100010000000	Align1 - xyzw ,(8),
20	0000100010000010	Align1 NoMask - xyzw ,(1),
21	0011000010000000	Align1   xyzw ,(8),
22	0011000000000001	Align16  ,(8),
23	0010000000000001	Align16  ,(4),
24	0011000000000010	Align1 NoMask  ,(8),
25	0011000000000101	Align16 NoDDClr  ,(8),
26	0011000000000101	Align16 NoDDChk  ,(8),
27	0011000000001000	Align1 SecHalf  ,(8),
28	0011000000000011	Align16 NoMask  ,(8),
29	0011000000000100	Align1 NoDDClr  ,(8),
30	0011000010000100	Align1 NoDDChk   xyzw ,(8),
31	0010000000000100	Align16 NoDDChk  ,(4),
Datatypes Bits		
0	0010011100000000	aud aud id <1> dir
1	0010001100001000	aud rud iud <1> dir
2	0010011100000000	rud aud id <1> dir
3	0010000000110000	aud iud aud <1> dir
4	0010101101001010	ruw ruw iuw <1> dir
5	0010000001101011	rw rw aud <1> dir
6	0011000110001011	aw rub rub <1> dir
7	0010111101101011	rw rw iww <1> dir



<b>Control Bits</b>		
8	001000000111101100	aw iw aud <1> dir
9	001000000001100001	rud jud aud <1> dir
10	001000110010100101	rd rd jud <1> dir
11	001000000001000001	rud mud aud <1> dir
12	001000001000110001	rub rub aud <1> dir
13	001000001000101001	ruw rub aud <1> dir
14	001000000000100000	aud rud aud <1> dir
15	001000001000110010	mub rub aud <1> dir
16	001010010100101001	ruw ruw ruw <1> dir
17	001011010010100101	rd rd rw <1> dir
18	001000000110100101	rd rw aud <1> dir
19	001100011000101001	ruw rub rub <1> dir
20	001011011000101100	aw rub rw <1> dir
21	001011010110100101	rd rw rw <1> dir
22	001011110110100101	rd rw iw <1> dir
23	001111011110111101	rf rf rf <1> dir
24	001111011110111100	af rf rf <1> dir
25	001111011110111101	rf rf rf dir
26	001111011110011101	rf af rf <1> dir
27	001111011110111110	mf rf rf <1> dir
28	001000000000100001	rud rud aud <1> dir
29	001000000000100010	mud rud aud <1> dir
30	00100111111011101	rf mf id <1> dir
31	001000001110111110	mf rf aud <1> dir
<b>Source Operands Bits</b>		
0	000000000000	dir <0;1,0>
1	010110001000	dir <16;16,1>
2	010001101000	dir <8;8,1>
3	001000101000	dir <2;2,1>
4	011010010000	dir <32;16,2>
5	000100100000	dir <1;2,0>
6	010001101100	indr <8;8,1>
7	010101110000	dir <16;8,2>





Control Bits		
8	011001111000	dir <32;8,4>
9	001100101000	dir <4;2,1>
10	010110001100	indr <16;16,1>
11	001000100000	dir <2;2,0>
12	010110001010	- dir <16;16,1>
13	000000000010	- dir <0;1,0>
14	010101010000	dir <16;4,2>
15	010101101000	dir <16;8,1>
16	111101001100	indr <0;4,1>
17	111100101100	indr <0;2,1>
18	011001110000	dir <32;8,2>
19	010110001001	(abs) dir <16;16,1>
20	010101011000	dir <16;4,4>
21	001101001000	dir <4;4,1>
22	010000101100	indr <8;2,1>
23	010000000000	dir <8;1,0>
24	001101110000	dir <4;8,2>
25	001100010000	dir <4;1,2>
26	001100000000	dir <4;1,0>
27	010001101010	- dir <8;8,1>
28	001101111000	dir <4;8,4>
29	000001110000	dir <0;8,2>
30	001100100000	dir <4;2,0>
31	001101010000	dir <4;4,2>
Subregister Bits		
0	0000000000000000	0 0 0
1	000000000000100	4 0 0
2	000000110000000	0 12 0
3	111000000000000	0 0 28
4	011110000001000	8 0 15
5	000010000000000	0 0 1
6	000000000010000	16 0 0
7	000110000001100	12 0 3



Control Bits		
8	0010000000000000	0 0 4
9	0000010000000000	0 16 0
10	000001010010100	20 20 0
11	000000001010110	22 2 0
12	0100000000000000	0 0 8
13	1100000000000000	0 0 24
14	0001000000000000	0 0 2
15	0000000100000000	0 4 0
16	000000000001000	8 0 0
17	1000000000000000	0 0 16
18	0000010100000000	0 20 0
19	0010100000000000	0 0 5
20	0011000000000000	0 0 6
21	000000001010100	20 2 0
22	101101010010100	20 20 22
23	0101000000000000	0 0 10
24	000000010001111	15 4 0
25	0110000000000000	0 0 12
26	1111100000000000	0 0 31
27	1010000000000000	0 0 20
28	000000000001111	15 0 0
29	000100010001111	15 4 2
30	001000010001111	15 4 4
31	0001100000000000	0 0 3

## 5.5.2 Hardware Behavior

Upon receiving an instruction with the bit[29] CompactCtrl bit set, HW uses the 5 indexes inside the compacted instructions to lookup the compaction table, then uses the table output to reconstruct the full size instruction.

If any source of the compacted instruction is immediate, only 13bits of the immediate value is encoded in the compacted instruction, HW sign extends bit[12] all the way for the entire immediate DWord.



**Table 5-26. DevSNB Compacted Instruction Format**

DW #	Instr Bits Alloc	High Bit	Low Bit	Instr Bits Used	Description	Bits in 128bits Format	Description (Imm. Src0 or Src1)	Bits in 128bits Format (Imm. Src0 or Src1)
1	8	63	56	8	Src1 RegNum	[108:101]	Imm[23:16] Imm[7:0]	[119:112] [103:96]
	8	55	48	8	Src0 RegNum	[76:69]	Src0 RegNum	[76:69]
	8	47	40	8	Dst RegNum	[60:53]	Dst RegNum	[60:53]
	5	39	35	5	Src1Index[4:0]	[120:109]	Src1Index[4:0]	[127:120] [111:104]
	3	34	32	3				
0	2	31	30	2	Src0Index[4:0]	[88:77]	Src0Index[4:0]	[88:77]
	1	29	29	1	CmpCtrl	[29]	CmpCtrl	[29]
	1	28	28	1	FlagSubRegNum	[89]	FlagSubRegNum	[89]
	4	27	24	4	CondModifier	[27:24]	CondModifier	[27:24]
	1	23	23	1	AccWrCtrl	[28]	AccWrCtrl	[28]
	5	22	18	5	SubRegIndex[4:0]	[100:96] [68:64] [52:48]	SubRegIndex[4:0]	[100:96] [68:64] [52:48]
	5	17	13	5	DataTypeInfo[4:0]	[63:61] [46:32]	DataTypeInfo[4:0]	[63:61] [46:32]
	5	12	8	5	ControlIndex[4:0]	[31] [23:8]	ControlIndex[4:0]	[31] [23:8]
	1	7	7	1	DebugCtrl	[30]	DebugCtrl	[30]
	7	6	0	7	Opcode	[6:0]	Opcode	[6:0]

**Definitions of Fields in the Compact Instruction**

Bits	Description
63:56	<b>Bits [108:101] Source1 register number</b> forms bits [108:101], the source 1 register number field. If immediate source is used, this field forms [103:96] of the 128-bit instruction word.
55:48	<b>Bits [76:69] Source0 register number</b> This field, after unpacking, forms bits [76:69], the source 0 register number field, of the 128-bit instruction word.
47:40	<b>Bits [60:53] Destination register number</b> This field, after unpacking, forms bits [60:53], the destination register number field, of the 128-bit instruction word.
39:35	<b>Src1Index</b> The 5-bit index for source 1. The 12-bit table-look-up result forms bits [120:109], the source 1 register region fields, of the 128-bit instruction word if immediate source is used, this field forms [108:104] of the 128-bit instruction word. Bit[39] is replicated to [127:109] of the 128-bit instruction word.
34:30	<b>Src0Index</b> The 5-bit index for source 0. The 12-bit table-look-up result forms bits [88:77], the source 0 register region fields, of the 128-bit instruction word.
29	<b>CompactCtrl – Compaction Control</b> This field indicates whether the instruction is in the 64-bit compaction form. When this bit is set (bit 29 of DW0), the instruction length is only 64-bit. The bit location is fixed in both 128-bit and 64-bit instruction forms. 0 = 128-bit form (normal) 1 = 64-bit compaction form
28	<b>Bit [89] – FlagSubRegNum</b> This field, after unpacking, is bit [89] of the 128-bit instruction word.
27:24	<b>Bits [27:24] – CondModifier</b> This field, after unpacking, is bits [27:24] of the 128-bit instruction word. The bit location is fixed in both 128-bit and 64-bit instruction forms.



Bits	Description
23	<b>AccWrCtrl – Implicit Accumulator Write Enable</b> This field, after unpacking, is bit[28] of the 128-bit instruction word.
22:18	<b>SubRegIndex</b> The 5-bit index for sub-register fields. The 15-bit table-look-up result forms bits [100:96], [68,64] and [52,48] of the 128-bit instruction word.
17:13	<b>DataTypeIdIndex</b> The 5-bit index for data type fields. The 18-bit table-look-up result forms bits [63:61] and [46, 32] of the 128-bit instruction word.
12:8	<b>ControlIndex</b> The 5-bit index for data type fields. The 17-bit table-look-up result forms bits[31], and [23, 8] of the 128-bit instruction word.
7	<b>Reserved</b>
6:0	<b>Bits [6:0] – Opcode</b> This field, after unpacking, is bits [6:0] of the 128-bit instruction word. The bit location is fixed in both 128-bit and 64-bit instruction forms.

### 5.5.3 Rules and Restrictions

In order to reduce the hardware complexity, the following rules and restrictions apply to the compressed instruction:

- Any branch/jump offset need to be based on the physical instruction size, after compaction.
- Any branch/jump instruction with immediate offset larger than 13bits should not be compacted.

### 5.6 End of Thread

There is no special instruction opcode (such as an END instruction) to cause the thread to terminate execution. Instead, the end of thread is signified by a *send* instruction with the end-of-thread (EOT) sideband bit set. Upon executing a *send* instruction with EOT set, the EU stops on the thread. Upon observing an EOT signal on the output message bus, the Thread Dispatcher makes the thread's resource available. If a thread uses pre-allocated resource managed by a fixed function, such as URB handles and scratch memory, some fixed function protocol also requires the thread to terminate with the message header phase to carry the information in order for the fixed function to release the pre-allocated resource.

EU hardware guarantees that if a terminated thread has in-flight read messages or loads at the time of 'end' that their writebacks will not interfere with either other threads in the system or new threads loaded in the system in the future.

More details can be found in the *send* instruction description in Instruction Reference chapter.



## 5.7 Creating Conditional Flags

FPU will output 2 sets of conditional signals, 1 set will be generated from before the adder outputs clamping/re-normalizing/format conversion logic, we call this the pre conditional signals. 1 set will be generated from the final results after clamping and re-normalizing/format conversion logic, and we will call this the post conditional signals. The post conditional signals are used for fusing the DX compare instruction. *The flags generated from the post conditional signals should be equivalent to the flags generated by a separate CMP instruction after the current arithmetic instruction.*

The pre conditional signals will be used to generate flags for CMP/CMPN instructions only, this logically does the compare of the 2 input sources. The post conditional signals will be used to generate flags for all the other arithmetic instructions, this logically does the compare of the result with zero.

CMPN with both sources are NaN is a don't care case since this doesn't impact the MIN/MAX operations.

The pre conditional signals include the following:

- **pre\_sign** bit: this bit reflects the sign of the computed result directly from the adders, without going through any kind of clamping, normalizing, or format conversion logic.
- **pre\_zero** bit: this bit reflects whether the value of the adder results are zero, again this should be obtained before any kind of clamping, normalizing, or format conversion logic.

The post conditional signals include the following:

- **post\_sign** bit: this bit reflects the sign of the final result after all the clamping, normalizing, or format conversion logic.
- **post\_zero** bit: this bit reflects whether the value of the adder results are zero after all the clamping, normalizing, or format conversion logic.
- **OF** bit: this bit reflects whether an overflow occurred in any of the computation of the current instruction, including clamping, re-normalizing, and format conversion.
- **NC** bit: The NaN computed bit indicates whether the computed result is not a number. It carries valid information for instructions operating on floating point values. For an operation on integer operands, this bit is always set to 0.
- **NS0** bit: The NaN bit indicates whether source 0 of an execution channel is not a number. It carries valid information for instructions operating on floating point values. For an operation on integer operands, this bit is always set to 0.
- **NS1** bit: The NaN bit indicates whether source 1 of an execution channel is not a number. It carries valid information for instructions operating on floating point values. For an operation on integer operands, this bit is always set to 0. For an operation with one source operand, this bit is also set to 0. This bit is only used for the comparison instruction CMPN, which is specifically provided to emulate MIN/MAX operations. For any other instructions, this bit is undefined.



**Flag Generation for CMP instructions** (The supported Conditional Modifiers are **.e**, **.ne**, **.g**, **.ge**, **.l**, and **.le**.)

Conditional Modifier	Meaning	Resulting Flag Value (for an execution channel)
'e'	Equal-to	<b>(pre_zero &amp; !(NS0   NS1))</b> . This conditional modifier tests whether the 2 sources are equal. If either source is NaN (i.e. NC is true), the flag is force to false.
'ne'	Not-Equal-to	<b>!(pre_zero &amp; !(NS0   NS1))</b> . This conditional modifier test whether the 2 sources are equal. It takes exactly the reverse polarity as modifier 'e'.
'g'	Greater-than	<b>(!pre_sign &amp; !pre_zero &amp; !(NS0   NS1))</b> . This conditional modifier tests whether source0 is greater than source1. If either source is a NaN (i.e. NC is true), the flag is forced to false.
'ge'	Greater-than-or-equal-to	<b>((!pre_sign   pre_zero) &amp; !(NS0   NS1))</b> . This conditional modifier tests whether source0 is greater than or equal to source1. If either source is a NaN (i.e. NC is true), the flag is forced to false.
'l'	Less-than	<b>(pre_sign &amp; !pre_zero &amp; !(NS0   NS1))</b> . This conditional modifier tests whether source0 is less than source1. If either source is a NaN (i.e. NC is true), the flag is forced to false.
'le'	Less-than-or-equal-to	<b>((pre_sign   pre_zero) &amp; !(NS0   NS1))</b> . This conditional modifier tests whether source0 is less than or equal to source1. If either source is a NaN (i.e. NC is true), the flag is forced to false.

**Flag Generation for CMPN instructions** (The supported Conditional Modifiers are **ge**, and **.l**) [DevSNB]

Conditional Modifier	Meaning	Resulting Flag Value (for an execution channel)
'ge'	Greater-than-or-equal-to	<b>(!pre_sign   pre_zero   (NS1 &amp; (Opcode==CMPN   Opcode==SELwCMod))) &amp; !(NS0 &amp; (Opcode==CMPN))</b> . This conditional modifier tests whether source0 is greater than or equal to source1. If source-1 is a NaN (i.e. NS is true), the flag is forced to true.
'l'	Less-than	<b>((pre_sign &amp; !pre_zero)   (NS1 &amp; (Opcode==CMPN   Opcode==SELwCMod))) &amp; !(NS0 &amp; (Opcode==CMPN))</b> . This conditional modifier tests whether source0 is less than source1. If source-1 is a NaN (i.e. NS is true), the flag is forced to true.



**Flag Generation for All Arithmetic Instructions other than CMP/CMPN instructions** (The supported Conditional Modifiers are **.e**, **.ne**, **.g**, **.ge**, **.l**, **.le**, **.r**, **.o**, and **.u**.)

Conditional Modifier	Meaning	Resulting Flag Value (for an execution channel)
<b>'e'</b>	Equal-to	<b>(post_zero &amp; !NC)</b> . This conditional modifier tests whether the 2 sources are equal. If either source is NaN (i.e. NC is true), the flag is forced to false.
<b>'ne'</b>	Not-Equal-to	<b>!(post_zero &amp; !NC)</b> . This conditional modifier test whether the 2 sources are equal. It takes exactly the reverse polarity as modifier <b>'e'</b> .
<b>'g'</b>	Greater-than	<b>(!post_sign &amp; !post_zero &amp; !NC)</b> . This conditional modifier tests whether source0 is greater than source1. If either source is a NaN (i.e. NC is true), the flag is forced to false.
<b>'ge'</b>	Greater-than-or-equal-to	<b>((!post_sign   post_zero) &amp; !NC)</b> . This conditional modifier tests whether source0 is greater than or equal to source1. If either source is a NaN (i.e. NC is true), the flag is forced to false.
<b>'l'</b>	Less-than	<b>(post_sign &amp; !post_zero &amp; !NC)</b> . This conditional modifier tests whether source0 is less than source1. If either source is a NaN (i.e. NC is true), the flag is forced to false.
<b>'le'</b>	Less-than-or-equal-to	<b>((post_sign   post_zero) &amp; !NC)</b> . This conditional modifier tests whether source0 is less than or equal to source1. If either source is a NaN (i.e. NC is true), the flag is forced to false.
<b>'o'</b>	Overflow	<b>(OF)</b> . This conditional modifier tests whether the computed result causes overflow – the computed result is outside the range of the destination data type. All other internal conditional signals are ignored.
<b>'u'</b>	Unordered	<b>(NC)</b> . This conditional modifier tests whether the computed result is a NaN (unordered). All other internal conditional signals are ignored.

## 5.8 Destination Hazard

The architecture has built-in hardware to avoid destination hazard.

Destination Hazard stands for the risk condition when multiple operations are trying to write to the same destination and the result of the destination may be ambiguous. This may or may not happen on DevSNB for two instructions with the same destination, or with destinations that have overlapped register region, depending on the ordering of the arrival of destination results. Let's consider two instructions in a thread with potential destination hazard. There may be other instruction between them as long as there is no instruction sourcing the same destination. Using register scoreboards, the hardware automatically takes care of the destination hazard by not issuing the second instruction until the destination scoreboard is cleared. However, for certain cases, in fact for most cases, such destination hazard indicated by the register scoreboard is false, causing unnecessary delay of instruction issuing. This may result in lower



performance. The destination dependency control field in the instruction word  $\{NoDDClr, NoDDhk\}$  allows software to selectively override such hardware destination dependency mechanism. Such performance optimization hooks must be used with extreme caution. When it is not 100% certainty that it is a false destination hazard, programmer should rely on hardware to resolve the dependency.

As the destination dependency control field does not apply to *send* instruction, there is only one condition that a programmer may use the  $\{NoDDClr, NoDDChk\}$  capability.

- If none of the two instructions is *send*, there CANNOT be any destination hazard. This is because instructions within a thread are dispatched in order (single-issued) and the execution pipeline is in-order and has a fixed latency.

## 5.9 Non-present Operands

Some instructions do not have two source operands and one destination operand. If an operand is not present for an instruction the operand field in the binary instruction must be filled with null. Otherwise, results are unpredictable.

Specifically, for instructions with a single source, it only uses the first source operand  $\langle src0 \rangle$ . In this case, the second source operand  $\langle src1 \rangle$  must be set to null and also with the same type as the first source operand  $\langle src0 \rangle$ . It is a special case when  $\langle src0 \rangle$  is an immediate, as an immediate  $\langle src0 \rangle$  uses DW3 of the instruction word, which is normally used by  $\langle src1 \rangle$ . In this case,  $\langle src1 \rangle$  must be programmed with register file ARF and the same data type as  $\langle src0 \rangle$ .

## 5.10 Instruction Prefetch

Due to prefetch of the instruction stream, the EUs may attempt to access up to 8 instructions (128 bytes) beyond the end of the kernel program – possibly into the next memory page. Although these instructions will not be executed, software must account for the prefetch in order to avoid invalid page access faults. One possible (though inefficient) solution would be to pad the end of all kernel programs with 8 NOOP instructions. A more efficient approach would be to ensure that the page after all kernel programs is at least valid (even if mapped to a dummy page). Note that the **General State Access Upper Bound** field of the STATE\_BASE\_ADDRESS command can be used to prevent memory accesses past the end of the General State heap (where kernel programs must reside).





## 6. Exceptions

### 6.1 Introduction

The architecture defines a basic exception handling mechanism for several exception cases. This mechanism supports both normal operations such as extensions of the mask-stack depth, as well as illegal conditions.

The following exception-types are supported:

Type	Trigger / Source	Sync/Async Recognition
MaskStack Overflow / Underflow	Hardware	Synchronous (w/ special case for 'do')
Software Exception	Thread code	Synchronous
Breakpoint	A bit in the instruction word	Synchronous
Illegal Opcode	Hardware	Synchronous
Halt	MMIO register write	Asynchronous

Threads may choose which exceptions to recognize and which to ignore. This mask information is specified on a per-kernel basis in fixed function state generated by the driver, and delivered to an EU as part of a new-thread dispatch. Upon arrival at the EU, the exception-mask information is used to initialize the exception enable fields of that thread's CR0.1 register, which controls exception recognition. This register is instantiated on a per-thread basis, allowing independent control of exception-type recognition across hardware threads. The exception enables in the CR0.1 register are r/w, and thus can be enabled/disabled via software at anytime during thread execution.

The exception handling mechanism relies on the "system routine", a single subroutine which provides common exception handling for all threads on all EUs in the system. This system routine is defined per-context and is identified via a 32b System-IP (SIP) register in context state. At the time of each context switch, the appropriate SIP for that context is loaded into each EU, allowing each context to have custom implementation of exception handling routines if so desired.

### 6.2 Exception-Related Architectural Registers

Exception-related registers are defined in architectural register CR0.0 through CR0.2. These registers are instantiated on a per-thread basis providing each hardware thread with unique control over exception recognition and handling. The registers provide the capability to mask exception types, determine the type of raised exception, provide storage the return address, and control exiting from the system routine back to the application thread.

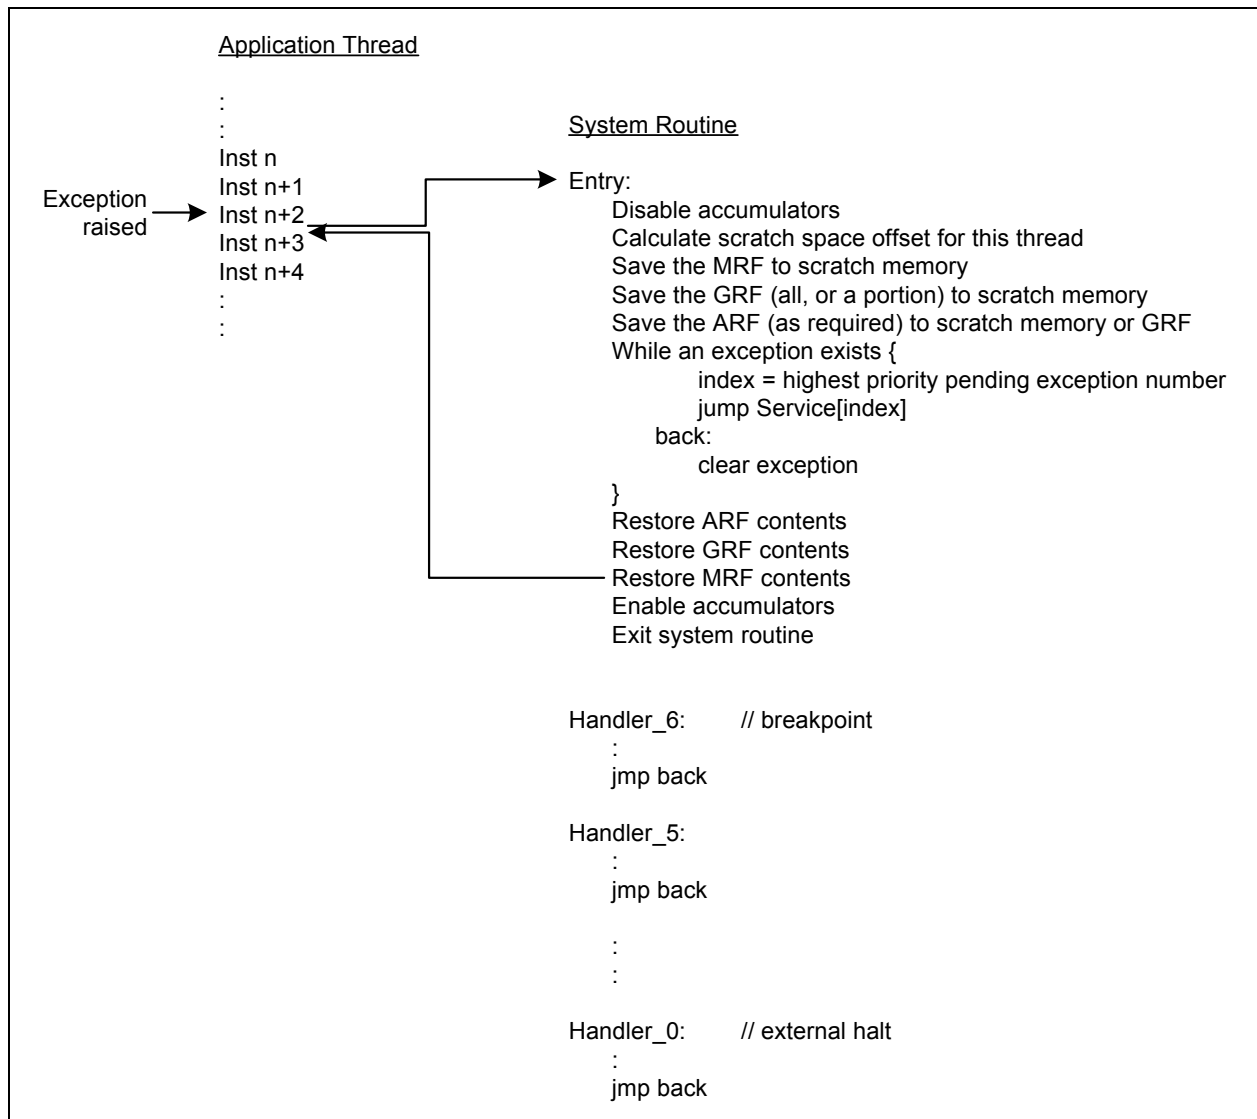


Many of the bits in these registers are manipulated by both hardware and software. In all cases, the read/write operations by hardware and software occur at exclusive times in a thread's lifetime, thus there is no need for an atomic R-M-W operation when accessing these registers.

## 6.3 System Routine

### 6.3.1 General Flow of the System Routine

The following diagram illustrates the basic flow of exception handling and structure of the system routine.





## 6.3.2 Invoking the System Routine

The system routine is invoked in response to a raised exception. Once an exception is raised, no further instructions from the application thread will be issued until the system routine has executed and returned control back to the application thread.

After an exception is recognized by hardware, the EU saves the thread's IP into its AIP register (CR0.2), and then moves the system routine offset, SIP, into the thread's IP register. At this point the next instruction to issue for that thread will be the first instruction of the system routine.

The system routine maintains the same execution priority, GRF and MRF register space, and thread state as that of the application thread from which it was invoked. Due to the assumption of the same priority, there may be significant absolute time between exception being raised and the actual invocation of the system routine, as other higher priority threads within the EU continue to execute. From a thread's perspective, once an exception is recognized, the next instruction to issue is from the system routine.

At the time of system routine invocation, there may still be outstanding registers in-flight from the application thread. Depending on the instruction sequence in the system routine, an in-flight register may be referenced by the system routine and cause a register-in-flight dependency. These dependencies are honored by the system routine and may cause the system routine to be suspended until such time that the register retires.

Exception processing is non-nested within a system routine. If a future exception is detected while executing the system routine, the exception is latched into CR0.1, but does not cause a nested re-invocation of the system routine. The exception recognition hardware recognizes only one outstanding exception of each type; i.e. once a specific exception type is detected and latched in CR0.1, and until the exception is cleared, any further exception of that type will be lost.

Accumulators are not natively preserved across the system routine. To make sure the accumulators are in the identical state once control is returned to the application thread, the system routine must either set the Accumulator Disable bit of CR0.0 prior to using any instruction which modifies an accumulator, or manually save/restore the accumulators (to GRF registers or system thread scratch memory) around the system routine. Saving/restoring accumulators, including their extended precision bits, can be accomplished by a short series of mov's and shifts of the accumulator register. Also note the state of the Accumulator Disable bit itself must be preserved unless, by convention, the driver software limits its manipulation to only the system routine.

Further, upon system routine entry, the execution-related masks (Continue, Loop, If, and Active masks, contained in the Mask Register) will remain set as they were in the application thread. Thus only a subset of channels may be active for execution. To enable execution on all channels, the system routine may choose to use the instruction option 'NoMask', or may choose to set the mask registers to the desired value so long as it saves/restores the original masks upon system routine entry/exit.

Similarly there is no hardware mechanism to preserve flags, mask-stacks, or other architectural registers across the system routine. The system routine must ensure that these values are preserved (see Section 6.3.7 for related discussion).

## 6.3.3 Returning to the Application Thread

Prior to returning control to the application thread, the system routine should clear the proper Exception Status and Control bit in CR1. Failure to do so will force the thread's execution to re-enter the system routine prior to any further instructions being executed from that application thread. (Note that single-stepping functionality is the one exception where the exception's Status and Control bit is not reset prior to exit.)



The system routine may choose to loop under a single invocation of the system routine until all pending exceptions are serviced, or may choose to service exceptions one at a time (a simpler solution, but less efficient).

The system routine is exited, and control returned to the application thread, via a write to the Master Exception State and Control bit of CR0.0. Upon clearing this bit, the value of the AIP architectural register (CR0.2) is restored to the thread's IP register and, with no further exceptions pending, execution resumes that address. The system routine must follow any write to Master Exception State and Control bit with at least one `simd-16 'nop'` instruction to allow control to transition. Throughout the system routine, the AIP register maintains its value at the time the exception was raised unless directly modified by the system routine. (See the AIP register definition for specifics on the IP value saved to AIP).

### 6.3.4 System-IP (SIP)

The System IP (SIP) is a 16B-aligned 32b offset of the first instruction of the system routine, relative to the General State Base Address. It is set via the `STATE_IP` command to the command streamer. The upper 28b of the 32b address is automatically delivered to all EUs.

When the system routine is invoked, the application thread's current IP is first saved into the AIP field of the thread's architectural register CR0.2. The SIP address is then loaded into the thread's IP register and execution continues within the system routine. Thus each invocation of the system routine has a common entry point at the first instruction of the system routine. Upon system routine completion, the value held in AIP is returned to IP and execution continues on the application thread at the place where the exception was recognized.

### 6.3.5 System Routine Register Space

The system routine uses the same GRF and MRF space at the thread which invoked it. As such all of the calling thread's registers and their contents are visible to the system routine. Further, the system routine must only use `r0..r15` of the GRF, as a minimal thread may have requested and been allocated this few. If the system routine requires more registers than this, the driver should establish a higher minimum allocation to all threads. It should also be noted that the system routine may encounter any residual register dependencies of the calling thread until such time that they clear by the return of in-flight writebacks.

Only one 32b GRF location, R0.4, is reserved for system routine usage. This is sufficient to allow the system routine to calculate the appropriate offset of its private scratch memory in the larger system-scratch memory space (as dictated by binding table entry 254). The offset is left as a driver convention, but likely based on a combination of Thread and EU IDs (see example system handler in section 6.3.6). Other than the reserved R0.4 register field, there is no explicit GRF register space dedicated to the system routine, and any GRF needs must be accomplished via: (a) convention between the system routine an application thread, or (b) the system routine temporarily spilling the thread's GRF register contents to scratch memory, and restoration prior to system routine exit.

No persistent storage is natively allocated to the system routine, although a driver implementation may choose to carve out a piece of system scratch memory though it own convention.

Any parameter passing to the system routine (for use by s/w exceptions) is performed via the GRF based on a system-thread/application-thread convention.



### 6.3.6 System-Scratch Memory Space

There is a single unified system-scratch memory space per context shared by all EUs. It is anticipated that block is further partitioned into a unique scratch sub-space per-thread via convention implemented in the system routine, with a each hardware thread having a uniform block size at a calculated offset from the base address. The block address for a thread is based on an offset derived from the thread's execution unit ID and thread ID made available through the TID and EUID field of architectural register SR0.0.

$$\text{Per\_Thread\_Block\_Size} = \text{System\_Scratch\_Block\_Size} / (\text{EU\_Count} * \text{Thread\_Per\_EU});$$

$$\text{Offset} = (\text{SR0.0.EID} * \text{Threads\_Per\_EU} + \text{SR0.0.TID}) * \text{Per\_Thread\_Block\_Size};$$

where in DevSNB...

$$\text{Threads\_Per\_EU} = 4$$

$$\text{EU\_Count} = 8$$

System\_Scratch\_Block\_Size is a driver choice

Access to the system-scratch memory is performed through the Data Port via linear single-register or block-based read/write messages. The driver may choose to use any binding table index for system-scratch surface description. As a practical matter, the same index is expected to be used across all binding tables, as the index is typically hard coded in dataport messages used within the system routine coupled with the fact that a single system instance routine is used for all threads. Read/write messages to the Data Port contain the address of the binding table (provided in R0 of all threads) and an offset, from which the Data Port calculates the final target address.

The size of the overall system-scratch memory block is a function of the system routine's feature set, but should be, at a minimum, of size sufficient to cover storing the entire GRF, MRF, flags, and architectural register set for all hardware threads. Additional per-thread memory may also be required for mask-stack spill/fill, system routine constants, and any global persistent storage needs of the system routine.

It is expected that the system-memory block is allocated by the driver at context-create time and remains persistent at a constant memory address throughout the context's lifetime.

### 6.3.7 Conditional Instructions Within System Routines

It is expected that most, if not all, control flow within the system routine is scalar in nature. If so, the system routine should set SPF (Single Program Flow, CR0.0) to enable scalar branching. In this mode, conditional/loop instructions do not update the mask-stacks and therefore do not have restrictions on their use nor require the save/restore of hardware mask-stack registers.

If SIMD branching is desired within the system routine, special considerations must be taken. Upon entry to the system routine, the depth of the mask-stacks is unknown at that point, and may be near-full. If so, a subsequent conditional instruction and its associated mask 'push' may cause a stack overflow. This would generate an exception-within-the-system-routine, an unsupported occurrence. To prevent this, if the system routine uses SIMD conditional instructions, it must save the mask-stacks prior to the first SIMD conditional instruction, and restore them after the last SIMD conditional instruction. As a general solution, it may be easiest simply to implement the save/restore as part of the entry/exit code sequence,



using an available GRF register-pair as storage location. Once saved, the stacks should be reset to their empty condition, namely depth = 0 and top-of-stack = 0xFFFFFFFF.

### 6.3.8 Messages in System Routines

The system routine uses the same MRF space as the thread on whose behalf the system routine was invoked. To allow the thread to resume with the same state as prior to the system routine invocation, the thread's MRF contents must be preserved across a system routine invocation. If the system routine requires MRF space for messages, it must manually save and restore the MRF locations which it uses.

Note that the MRF can only be used as an instruction's destination register, not a source. Therefore there is no option to save the MRF to the GRF. Thus the system routine should save the MRF contents to its dedicated scratch space. By convention it is recommended that MRF register m0 be reserved for system-thread use. This allows the system routine enough space to construct an initial Data Port write message starting at m0 without corrupting any MRF registers, facilitating a complete save/restore of the MRF by the system-thread.

### 6.3.9 Use of 'NoDDClr'

The instruction word defines an instruction option 'NoDDClr' which overrides the native register dependency clearing mechanism of the typical instruction. When specified, 'NoDDClr' does not clear, at register writeback time, the dependency placed on the destination register of the instruction. Use of this mechanism may provide increased performance when the kernel can guarantee no dependency issues between instructions, but may cause issues with exception handling in some circumstances as discussed here.

Typically 'NoDDClr' is used in an instruction series to enable a sequence of writes to sub-fields of a GRF register without paying a dependency penalty on each instruction. In this case, 'NoDDClr' and 'NoDDChk' are used across an instruction sequence to allow the first instruction to set the destination dependency, interior instructions to write to the GRF register w/o dependency checks, and the last instruction clear the dependency. (This sequence is referred to as a 'NoDDClr' code block going forward). By only allowing the last instruction to clear the dependency, program execution is prevented from going beyond a certain point until all writes of that sequence are known to retire.

The problem arises should an exception be raised within a 'NoDDClr' code block. In this case, there exists the potential for the system routine to hang while attempting to save/restore the code blocks destination register, as the outstanding dependency on that register will not clear until the final instruction of the block is executed – sometime after the system thread returns. Should the system routine attempt to use that register, the system routine will hang waiting on a dependency to clear from an instruction which has not yet been issued.

**This is a known condition and will in some cases not allow the full GRF contents to be externally visible in system routine scratch space during a break or halt exception.** To minimize the number of cases of such, guidelines are provided below for consideration. (Note that these are general guidelines, some of which can be alleviated through careful coding and register usage conventions and restrictions.)

- 'NoDDClr' code blocks should only be used where absolutely necessary.
- Instructions which may generate exceptions should not be placed within 'NoDDClr' blocks.
- If possible, use 'NoDDClr' on registers high in the thread's register allocation (e.g. r120), thus even if a system routine hang occurs, as much of the GRF is visible as possible. (Note this would also require the system routine to update the progress of the GRF dump, perhaps with each GRF block



written, or to initialize the system routine's scratch space to a known value, to be able to distinguish valid/locations from unwritten locations).

Also a driver implementation may consider a "disable-NoDDClr" option in which jitted code does not use the 'NoDDClr' capability. In this case, there is no change to the code that is jitted other than removal of the 'NoDDClr' instruction option. The code executes as normal, but with a higher number of thread switches in what would have been a NoDDClr code block.

## 6.4 Exception Descriptions

### 6.4.1 'Illegal' opcode

The ISA defines a single 'illegal' opcode. The byte value of the 'illegal' opcode is selected to be 0x00 due to it being a likely byte-value encountered by a wayward instruction pointer value. The 'illegal' instruction raises an exception prior to issue and operates as a 'nop' when issued down the execution pipeline. (Specifically, the opcode acts a 'nop', although other non-opcode instruction attributes still apply).

### 6.4.2 Undefined opcode

All undefined opcodes in the 8b opcode space are detected by hardware. If an undefined opcode is detected, the opcode is overridden by hardware, forcing it to the defined 'illegal' opcode. The offending instruction, should it eventually be issued down the execution unit's pipeline, generates an 'illegal opcode' exception as described in section 6.4.1. Note that the memory location of the offending opcode remains modified and may be queried if desired to determine its original value.

### 6.4.3 Software Exception

A mechanism is provided to allow an application thread to invoke an exception and is triggered through of the Software Exception Set and Clear bit of CR0.1. Sub-function determination and parameter passing into and out-of the exception handler is left to convention between the system-thread and application-thread. The thread's AIP instruction pointer is incremented prior to system-routine entry, therefore causing execution to resume at the subsequent application-thread instruction when the system routine is exited.

### 6.4.4 Breakpoint

A single-stepping capability may be implemented by leaving the "Breakpoint Exception Status and Control" set, and clearing the Breakpoint Suppress field prior to system routine exit. This combination causes the instruction associated with the breakpoint to be reissued, this time with the breakpoint suppressed, and then re-entry to the system routine prior to the subsequent instruction due to the lingering breakpoint exception that remained un-cleared.

### 6.4.5 External Halt

A 'halt' exception may occur via direction manipulation of a MMIO bit by driver software. The halt exception is sent to all EUs simultaneously (although no guarantee is made as to recognition in identical clocks). An EU recognizes this condition internally by generating an External Halt exception. A likely implementation of a handling routine would dump the thread's state to programmer-visible memory (such as the system routine's scratch space) for inspection purposes. Although generally recognized within a



few clocks, there is no specification as to the latency between triggering the Halt condition and it being recognized by an EU.

## 6.5 Events Which Do Not Generate Exceptions

The following conditions are either not recognized or do not generate an exception.

### Illegal Instruction Format

This includes malformed instructions in which the opcode is legal, but the source or destination operands, or instruction attributes are not compliant with the instruction specification. There is no direct hardware support to detect these cases and the outcome of issuing a malformed instruction is undefined. Note that DevSNB does not support self-modifying code, therefore the driver has an opportunity to detect such cases before the thread is placed in service.

### Malformed Message

A messages contents, destination registers, lengths, and descriptors are not interpreted in anyway by the execution units. Errors in specifying any of these fields do not raise exceptions in the execution unit but may be detected and reported by the shared functions.

### GRF Register Out-of-Bounds

Unique GRF storage is allocated to each thread which, at a minimum, satisfies that the register requirements specified in the thread's declaration. References to GRF register numbers beyond that called for in the thread's declaration do not generate exceptions. Depending on implementation, out-of-bounds register numbers may be remapped to r0..r15, although this functionality should not be relied upon by the thread. The hardware guarantees the isolation of each threads register space, thus there is no possibility of direct register manipulation from an out-of-bounds register access.

### MRF Register Out-of-Bounds

A fixed amount of MRF register space is allocated for each thread, namely m0 through m23. References to MRF registers beyond m23 do not generate exceptions. Depending on implementation details, out-of-bounds register numbers may alias to in-bounds register numbers, although this functionality should not be relied upon by the thread.

### Hung Thread

There is no hardware mechanism in the execution units to detect a hung thread, and should it occur, the thread remains hung indefinitely. It is the expectation that one or more hung threads will eventually cause the driver to recognize a context timeout and take appropriate recovery action.

### Instruction Fetch Out of Bounds

The EUs implement a full 32b instruction address range (with the 4 lsb's don't care), making it possible for a thread to attempt to jump to any 16B aligned offset in the 32b address space. The EU itself does not provide any type of address checking on its instruction request stream sent to the memory/cache hierarchy, although various memory address related error conditions are reported through the Memory Interface Registers (specifically "Page Table Error Register").

### FPU Math Errors

The EU's floating point units have defined behavior for traditional floating point errors and do not generate exceptions. Therefore there is no support for signaling FPU math errors as exceptions.





## Destination Register Overflow

Depending on source operand contents, destination register size, and operation being performed, overflows may occur in the EU's pipeline. These are not flagged as exceptions and software must explicitly check the overflow bit in the thread's architectural register if overflow is a concern.

## 6.6 System Handler Example

The following code sequence illustrates some concepts of the system routine. It is intended to be just a shell, without getting into the specifics of each exception handler. The example frees enough MRF and GRF space to get the routine started, then jumps to the handler for the specific exception. Many other implementations are also valid, including single exception servicing (as opposed to looping) per invocation, and saving only the GRF or MRF space required by the exception being serviced.

```
#define ACC_DISABLE_MASK          0xFFFFFFFFD
#define MASTER_EXCP_MASK        0x7FFFFFFF
#define SYSROUTINE_SCRATCH_BLKSIZE 16384      //for example

// --- SharedFunc IDs ---
#define DPR                      0x04000000
#define DPW                      0x05000000

// --- message lengths ---
#define ML5                      0x00500000
#define ML9                      0x00900000

// --- response lengths ---
#define RL0                      0x00000000
#define RL4                      0x00040000
#define RL8                      0x00080000

// --- dataport block sizes ---
#define BS1_LOW                  0x0000
#define BS1_HIGH                 0x0100
#define BS2                      0x0200
#define BS4                      0x0300
```



```
// --- Scratch Layout ---  
#define SCR_OFFSET_MRF                0  
#define SCR_OFFSET_GRF                512           // + 16 reg  
#define SCR_OFFSET_ARF                512 + 4096   // + 16 + 128 reg
```

```
// --- Write Dataport constants ---  
// target=dcache, type= oword_block_wr, binding_tbl_offset=0  
#define DPW                            0x000
```

```
// --- Read Dataport constants ---  
// target=dcache, type= oword_block_rd, binding_tbl_offset=0  
#define DPR                            0x000
```

Sys\_Entry:

```
// --- calc scratch offset for this thread into r0.4 ---  
shr (1) r0.4 sr0.0:uw 6                {NoMask}  
add (1) r0.4 r0.4 sr0.0:ub            {NoMask}  
mul (1) r0.4 r0.4 SYSROUTINE_SCRATCH_BLKSIZE {NoMask}  
  
// --- setup m0 w/ block offset  
mov (8) m0 r0                          {NoMask}  
  
// --- save mrf 7...0; (may choose to save the whole mrf)  
add (1) m0.2 r0.4 SCR_OFFSET_MRF      {NoMask}  
send (8) null m0 null DPW|ML9|RL0     {NoMask}  
  
// --- save mrf 8...15; (optional; req'ed if sys-routine stays w/in mrf7-0)  
mov (8) m7 r0                          {NoMask}  
add (1) m7.2 r0.4 (SCR_OFFSET_MRF + 256) {NoMask}  
send (8) null m7 null DPW|ML9|RL0     {NoMask}
```



```
// --- save r0..r1 to system scratch ---  
  
// --- (Note: done as a single register to guarantee external  
// --- visibility — see “Use of ‘NoDDClr’” in Excpetions Bspec chapter  
  
mov (16) m1 r0 {NoMask}  
send (8) m0 null null DPW|ML2|RL0 {NoMask}  
  
// --- save r2..r3 to free some room  
  
mov (16) m3 r2 {NoMask}  
add (1) m0.2 r0.4 SCR_OFFSET_GRF + 64 {NoMask}  
send (8) m0 null null DPW|ML4|RL0 {NoMask}  
  
// --- save r4..r7 to free some room (optional, depending on needs)  
  
mov (16) m8 r4 {NoMask}  
mov (16) m10 r6 {NoMask}  
add (1) m7.2 r0.4 (SCR_OFFSET_GRF + 128) {NoMask}  
send (8) m7 null null DPW|ML5|RL0 {NoMask}  
  
// --- save r8..r11 to free some room (optional, depending on needs)  
  
mov (16) m1 r8 {NoMask}  
mov (16) m3 r10 {NoMask}  
add (1) m0.2 r0.4 (SCR_OFFSET_GRF + 256) {NoMask}  
send (8) m0 null null DPW|ML5|RL0 {NoMask}  
  
// --- save r12..r15 to free some room (optional, depending on needs)  
  
mov (16) m8 r12 {NoMask}  
mov (16) m10 r14 {NoMask}  
add (1) m7.2 r0.4 (SCR_OFFSET_GRF + 384) {NoMask}  
send (8) m7 null null DPW|ML5|RL0 {NoMask}  
  
// --- save ARF registers (optional, depending on use) ---  
// flags, maskstacks, others...
```



```
// --- save f0.0 ---  
mov (1) r1.0:uw f0.0 {NoMask}  
  
Next: // --- exceptions pending? If not, exit ---  
cmp.e (1) f0.0 cr0.4:uw 0:uw {NoMask}  
(f0.0) mov (1) IP EXIT {NoMask}  
  
// --- find highest priority exception ---  
lzd (1) r1.1:uw cr0.4:uw {NoMask}  
  
// --- jumtable to service routine ---  
jmp (1) r1.1:uw {NoMask}  
mov (1) IP CRService_0 {NoMask}  
mov (1) IP CRService_1 {NoMask}  
mov (1) IP CRService_2 {NoMask}  
// :  
// :  
// :  
mov (1) IP CRService_15 {NoMask}  
  
mov (1) IP Next  
  
Service_0:  
// clear exception from CR0.1  
// perform service routine  
// jump to exit (or if looping on exceptions, go to next loop)  
  
// :  
// :  
  
Service_15:  
// clear exception from CR0.1
```



```
// perform service routine  
// jump to exit (or if looping on exceptions, go to next loop)
```

Exit:

```
// --- restore f0.0 ---  
  
// --- restore ARF registers (as required) ---  
// flags, maskstacks, others...  
  
// --- restore r12..r15 ---  
// --- restore r8..r11 ---  
// --- restore r4..r7 ---  
// --- restore r0..r3 ---  
  
// --- restore m8..m15 ---  
// --- restore m0..m7 ---  
  
// --- clear Master Exception State bit in CR0.0  
and (1) cr0.0 cr0.0 MASTER_EXCP_MASK  
nop (16)
```

Below is a code sequence to programmatically clear the GRF scoreboard in the case of a timeout waiting on a register that may never return:



```
// At this point, all we know is we have a hung thread. To get around
// any hung dependency, we can walk the GRF using NoDDChk, using execution mask
// of f0 = 0 so we don't touch the register contents.
```

Clear\_Dep:

```
    mov f0 0x00
(f0) mov r0 0x00 {NoDDChk}
(f0) mov r1 0x00 {NoDDChk}
(f0) mov r2 0x00 {NoDDChk}
    ...
    ...
(f0) mov r127 0x00 {NoDDChk}

// GRF scoreboard now cleared.
```



# 7. Instruction Set Summary

## 7.1 Instruction Set Characteristics

### 7.1.1 SIMD Instructions and SIMD Width

Instructions are SIMD (single instruction multiple data) instructions. The number of data elements per instruction, or the execution size, depends on the data type. For example, the execution size for instructions operating on 256-bit wide vectors can be up to 8 for 32-bit data types, and be up to 16 for 16-bit data. The maximum execution size for instructions for 8-bit data types is also limited to 16.

An instruction compression mode is supported for 32-bit instructions (including mixed 32-bit and 16-bit data computation). A compressed instruction works on twice as many SIMD data as that for a non-compressed instruction. Non-compressed instructions are also referred to as 'native' instructions. A compressed instruction is converted into two native instructions by the instruction dispatcher in the EU.

Instructions are executed on a narrower SIMD execution pipeline. Therefore, native instructions take multiple execution cycles to complete. See Table 7-1 for parameters for difference device hardware.

**Table 7-1. GRF instruction execution parameters in device hardware**

Device	Execution Pipeline	Native Instruction	
	Width (bits)	Max Width (bits)	Min Execution Cycles (clocks)
[DevSNB]	128	512	4

### 7.1.2 Instruction Operands and Register Regions

The majority of instructions may have up to three operands, two sources and one destination. Each operand is able to address a register region. Source operands support negate and absolute modifier and channel swizzle, and the destination operand supports channel mask.

Dual destination instructions are also supported (four-operand instructions in a general sense): One case is for the implied destination – flag register, where the conditional modifiers and the predicate modifiers may apply. Another case is the message header creation (implied move or implied assembling of the header) in the *send* instruction.

Each execution channel contains an accumulator that is wider than the input data to support back-to-back accumulation operations with increased precision. The added precision (see accumulator register description in Execution Environment chapter) determines the maximum number of accumulations before possible overflow. The accumulator can be pre-loaded through the use of *mov*. It can also be pre-loaded by arithmetic instructions such as *add*, *mul*, since the result of these instructions can go to the accumulator. The accumulator registers are per thread and therefore safe for thread switching.

Register access can be direct or register-indirect. Register-indirect register access uses address registers plus an immediate offset term to compute the register addresses, and only applies to the first source operand (src0) and/or the destination operand.



There is one address register that contains 8 sub-registers. Each sub-register contains a 16-bit unsigned value. The leading two sub-registers form a special doubleword that can be used as the descriptor for the send instruction.

Source operand can also be immediate value (also referred to as inline constants). For instructions with two source operands, only the second operand <src1> is allowed to be immediate. For instructions with only one source operand, the source operand <src0> is used and it can be an immediate.

An immediate source operand can be a scalar value of specified type up to 32-bit wide, which is replicated to create a vector with length of Execution Size. An immediate operand can also be a special 32-bit vector with 8 elements each of 4-bit signed integer value, or a 32-bit vector with 4 elements each of 8-bit restricted float value.

### 7.1.3 Instruction Execution

It is implied that all instructions operate across all channels of data unless otherwise specified either via destination mask, predication, execution mask (caused by SIMD branch and loop instructions), or execution size.

Instruction execution size can be specified per instruction, from scalar (*ExecSize* = 1) up to the maximal execution size supported for the data type, with the restriction that execution size can only be in power of 2.

## 7.2 Instruction Machine Formats

This section shows the machine formats of the instruction set. The instructions in the architecture have fixed length of 128 bits. Out of the 128 bits, there are 120 bits in use, and the remaining bits are reserved for future extensions. One instruction consists of instruction fields that control various stages of execution of the instruction. These fields are roughly groups into the 4 doublewords as the following.

- Instruction Operation Doubleword (DW0) contains the opcode and other general instruction control fields.
- Instruction Destination Doubleword (DW1) contains the destination operand (<dst>) and the register file and type of source operands.
- Instruction Source-0 Doubleword (DW2) contains the first source operand (<src0>) and flag register number
- Instruction Source-1 Doubleword (DW3) contains the second source operand (<src1>) and is used to hold the 32-bit immediate source (imm32 as <src0> or <src1>).

Figure 5-1. DevSNB Instruction Format – 1-src and 2-src depicts the details of the organization of fields in the 128-bit instruction word based on the Addressing Mode and Access Mode of an instruction. Definitions for individual instruction fields are provided in the following sections.

The **send** instruction is shown in the table as it has some unique instruction fields. For example, the message descriptor (plus EOT) occupies the whole DW3, and the immediate destination register overlaps with the Conditional Modifier field. The rest of fields in DW0-3 follows the definition on the left, depending on Addressing Mode and Access Mode of the **send** instruction.

The **math** and **conditional branch** instruction are also shown in the table as they have some unique instruction fields.





Not shown is for immediate operands. When an immediate source is present in an instruction, it always occupies the whole DW3 with a 32-bit value.

Support for indirect addressing for <src1>, as shown by the gray areas in Table 4-2 is device dependent. See Table 5-25 (Indirect source addressing support available in device hardware) in ISA Execution Environment for details.

**Figure 5-1. DevSNB Instruction Format – 1-src and 2-src SNB**

DW #	Instr Bits Alloc	High Bit	Low Bit	Instr Bits Used	AddrMode = Direct		AddrMode = Indirect		SEND		MATH	Conditional Branch	Break Continue	Imm Src	
					AccessMode = Align16	AccessMode = Align1	AccessMode = Align16	AccessMode = Align1	MsgDesc Imm	MsgDesc Reg					
3	1	127	127	1	EOT										
	2	126	125	2											
	4	124	121	4											
	4	120	117	4											
	1	116	116	1	Src1.VertStride										
	2	115	114	2	Src1.Width		Src1.Width								
	2	113	112	2	Src1.ChanSel[7:4]	Src1.HorzStride	Src1.ChanSel[7:4]	Src1.HorzStride					UIP[15:0]		
	1	111	111	1	Src1.AddrMode										
	2	110	109	2	Src1.SrcMod										
	3	108	106	3	Src1.AddrSubRegNum										
	5	105	101	5	Src1.RegNum [7:0]										
	1	100	100	1	Src1.SubRegNum [4]	Src1.AddrImm [9:4]	Src1.ChanSel[3:0]	Src1.AddrImm [9:0]							
4	99	96	4	Src1.ChanSel[3:0]	Src1.SubRegNum [4:0]	Src1.ChanSel[3:0]	Src1.AddrImm [9:0]	Imm[28:0]	Reg32	Same	Same	JIP[15:0]	Imm[31:0]		
2	6	95	90	6	FlagSubRegNum										
	1	89	89	1											
	4	88	85	4	Src0.VertStride										
	1	84	84	1											
	2	83	82	2	Src0.Width		Src0.Width								
	2	81	80	2	Src0.ChanSel[7:4]	Src0.HorzStride	Src0.ChanSel[7:4]	Src0.HorzStride							
	1	79	79	1	Src0.AddrMode										
	2	78	77	2	Src0.SrcMod										
	3	76	74	3	Src0.AddrSubRegNum										
	5	73	69	5	Src0.RegNum [7:0]										
	1	68	68	1	Src0.SubRegNum [4]	Src0.AddrImm [9:4]	Src0.ChanSel[3:0]	Src0.AddrImm [9:0]							
	4	67	64	4	Src0.ChanSel[3:0]	Src0.SubRegNum [4:0]	Src0.ChanSel[3:0]	Src0.AddrImm [9:0]	Same	Same	Same	JIP[15:0]	Same	Same	
1	1	63	63	1	Dst.AddrMode										
	2	62	61	2	Dst.HorzStride		Dst.HorzStride								
	3	60	58	3	Dst.AddrSubRegNum										
	5	57	53	5	Dst.RegNum [7:0]										
	1	52	52	1	Dst.SubRegNum [4]	Dst.AddrImm [9:4]	Dst.ChanEn[3:0]	Dst.AddrImm [9:0]							
	4	51	48	4	Dst.ChanEn[3:0]	Dst.SubRegNum [4:0]	Dst.ChanEn[3:0]	Dst.AddrImm [9:0]	Same	Same	Same	JIP[15:0]	Same	Same	
	1	47	47	0											
	3	46	44	3	Src1.SrcType										
	2	43	42	2	Src1.RegFile										
	3	41	39	3	Src0.SrcType										
	2	38	37	2	Src0.RegFile										
	3	36	34	3	Dst.DstType										
2	33	32	2	Dst.RegFile				Same	Same	Same	Same	Same	Same		
0	1	31	31	1	Saturate										
	1	30	30	1	DebugCtrl										
	1	29	29	1	CmpCtrl										
	1	28	28	1	AccWrCtrl				Same	Same		Same			
	4	27	24	4	CondModifier				SFID[3:0]	FC[3:0]		MBZ			
	3	23	21	3	ExecSize										
	1	20	20	1	PredInv										
	4	19	16	4	PredCtrl				Same	Same					
	2	15	14	2	ThreadCtrl					FC[5:4]					
	2	13	12	2	QtrCtrl										
	2	11	10	2	DepCtrl										
	1	9	9	1	WECtrl										
1	8	8	1	AccessMode				Same	Same	Same	Same	Same			
1	7	7	0	(reserved for Opcode)											
0	7	6	0	Opcode				Same	Same	Same	Same	Same			

DevSNB also supports a separate instruction format for instructions with 3 sources. It has fixed length of 128 bits. Out of the 128 bits, there are 114 bits in use, and the remaining bits are reserved for future extensions. One instruction consists of instruction fields that control various stages of execution of the instruction. These fields are roughly groups into the 4 doublewords as the following.

- Instruction Operation Doubleword (DW0) contains the opcode and other general instruction control fields.
- Instruction Destination Doubleword (DW1) contains the destination operand (<dst>), the register file and source modifiers, and flag subregister number.



- Instruction Source Doublewords (DW2/DW3) contains the three source operands (<src0>, <src1>, <src2>).

The 3-src instructions have the following restrictions compare to the 1-src/2-src instructions.

- The only supported instructions are: LRP, MAD, BFE, BFI2
- Only GRF register allowed for sources, and only GRF/MRF register allowed for destination
- Subregister number can only go down to DWord granularity.
- Must be Align16, uses Align16 style swizzling, with extra replication control. No other regioning support.



Figure 7-3 DevSNB Instruction Format – 3-src SNB

DW #	Instr Bits Alloc	High Bit	Low Bit	Instr Bits Used	Description
2,3	2	127	126	0	<i>reserved</i>
	8	125	118	8	Src2 Regnum
	3	117	115	3	Src2 Subregnum
	8	114	107	8	Src2 Swizzle
	1	106	106	1	Src2 RepCtrl
	1	105	105	0	<i>reserved</i>
	8	104	97	8	Src1 Regnum
	3	96	94	3	Src1 Subregnum
	8	93	86	8	Src1 Swizzle
	1	85	85	1	Src1 RepCtrl
	1	84	84	0	<i>reserved</i>
	8	83	76	8	Src0 Regnum
	3	75	73	3	Src0 Subregnum
	8	72	65	8	Src0 Swizzle
	1	64	64	1	Src0 RepCtrl
	1	8	63	56	8
3		55	53	3	Dst Subregnum
4		52	49	4	Dst chan enable
7		48	42	0	<i>reserved</i>
2		41	40	2	Src2 Modifier
2		39	38	2	Src1 Modifier
2		37	36	2	Src0 Modifier
2		35	34	0	<i>reserved</i>
1		33	33	1	FlagSubRegNum
1		32	32	1	Dst Regfile
0	1	31	31	1	<i>Saturate</i>
	1	30	30	1	<i>DebugCtrl</i>
	1	29	29	1	<i>CmptCtrl</i>
	1	28	28	1	<b>AccWrCtrl</b>
	4	27	24	4	<i>CondModifier</i>
	3	23	21	3	<i>ExecSize</i>
	1	20	20	1	<i>PredInv</i>
	4	19	16	4	<i>PredCtrl</i>
	2	15	14	2	<i>ThreadCtrl</i>
	2	13	12	2	<i>QtrCtrl</i>
	2	11	10	2	<i>DepCtrl</i>
	1	9	9	1	<b>WECtrl</b>
	1	8	8	1	<i>AccessMode</i>
	1	7	7	0	<i>(reserved for Opcode)</i>
7	6	0	7	<i>Opcode</i>	



## 7.2.1 Common Instruction Fields

### 7.2.1.1 1-src and 2-src Instructions

As shown in the table below, the meanings (encoding) of certain bit fields in the 128-bit instruction word varies depending on the values of other bit fields.

Table 7-2 provides the definition of common fields in the instruction word. The 'Width' column specifies the width of the field in bits. These common fields will be referred to later in describing the fields of different doublewords of the instruction. The definition for fields that have unique representation can be found in its corresponding doubleword of the instruction.

**Table 7-2. Definitions of Common Instruction Fields**

Field	Description	Width
<b>CondModifier</b>	<p><b>Conditional Modifier.</b> This field sets the flag register based on the internal conditional signals output from the execution pipe such as sign, zero, overflow and NaNs, etc. If this field is set to 0000, no flag registers are updated. Flag registers are not updated for instructions with embedded compares.</p> <p>This field may also be referred to as the <i>flag destination control</i> field.</p> <p>This field applies to all instructions except <i>send</i>, <i>sendc</i>, and <i>math</i>.</p> <p>0000 = Do not modify the flag register (normal)            0001 = Zero or Equal ('.z' or '.e')            0010 = Not Zero or Not Equal ('.nz' or '.ne')            0011 = Greater-than ('.g')            0100 = Greater-than-or-equal ('.ge')            0101 = Less-than ('.l')            0110 = Less-than-or-equal ('.le')            0111 = Reserved            1000 = Overflow ('.o')            1001 = Unordered with Computed NaN ('.u')            1010 -1111 = Reserved</p>	4
<b>AddrMode</b>	<p><b>Addressing Mode.</b> This field determines the addressing method of the operand. When it is cleared, the register address of the operand is directly provided by bits in the instruction word. It is called a direct register addressing mode. When it is set, the register address of the operand is computed based on the address register value and an address immediate field in the instruction word. This is referred to as a register-indirect register addressing mode.</p> <p>This field applies to the destination operand and the first source operand, &lt;src0&gt;. Support for &lt;src1&gt; is device dependent. See Table XX (Indirect source addressing support available in device hardware) in ISA Execution Environment for details.</p> <p>0 = "Direct". Direct register addressing            1 = "Register-Indirect" (or in short "Indirect"). Register-indirect register addressing</p>	1



Field	Description	Width
<p><b>RegNum</b></p>	<p><b>Register Number.</b> This field provides the register number for the operand. For GRF or MRF register operand, it provides the portion of register address aligning to 256-bit. For an ARF register operand, this field is encoded such that MSBs identify the architecture register type and LSBs provide its register number.</p> <p>This field together with the corresponding <i>SubRegNum</i> field provides the byte aligned address for the origin of the register region. Specifically, this field provides bits [12:5] of the byte address, while <i>SubRegNum</i> field provides bits [4:0].</p> <p>This field applies to the destination operand and the source operands. It is ignored (or not present in the instruction word) for an immediate source operand.</p> <p>This field is present if the operand is in direct addressing mode; it is not present if the operand is register-indirect addressed.</p> <p>Format = U8, if <i>RegFile</i> = <b>GRF</b>.</p> <p>0x00 to 0x7F = Register number in the range of [0, 127] 0x80 to 0xFF = Reserved</p> <p>Format = U8, if <i>RegFile</i> = <b>MRF</b>.</p> <p>0x00 to 0x0F = Register number in the range of [0, 15] 0x10 to 0xFF = Reserved</p> <p>Format = 8-bit encoding, if <i>RegFile</i> = <b>ARF</b>.</p> <p>This field is used to encode the architecture register as well as providing the register number. See Execution Environment chapter for details.</p>	<p>8</p>
<p><b>SubRegNum</b></p>	<p><b>Sub-Register Number.</b> This field provides the sub-register number for the operand. For GRF or MRF register operand, it provides the byte address within a 256-bit register. For an ARF register operand, this field also provides the sub-register number according to special encoding for the given architecture register.</p> <p>This field together with the corresponding <i>RegNum</i> field provides the byte aligned address for the origin of the register region. Specifically, this field provides bits [4:0] of the byte address, while <i>RegNum</i> field provides bits [12:5].</p> <p>This field applies to the destination operand and the source operands. It is ignored (or not present in the instruction word) for an immediate source operand.</p> <p>This field is present if the operand is in direct addressing mode; it is not present if the operand is register-indirect addressed.</p> <p>Format = U5, if <i>RegFile</i> = <b>GRF</b> or <b>MRF</b></p> <p>0x00 to 0x1F = Sub-Register number in the range of [0, 31]</p> <p>Format = 5-bit encoding, if <i>RegFile</i> = <b>ARF</b>.</p> <p>This field is used to encode the architecture register as well as providing the register number. See Execution Environment chapter for details.</p>	<p>5</p>
<p><b>AddrSubRegNum</b></p>	<p><b>Address Sub-Register Number.</b> This field provides the sub-register number for the address register. The address register contains 8 sub-registers. The size of each sub-register is one word. The address register contains the register address of the operand, when the operand is in register-indirect addressing mode.</p> <p>This field applies to the destination operand and the source operands. It is ignored (or not present in the instruction word) for an immediate source operand.</p> <p>This field is present if the operand is in register-indirect addressing mode; it is not present if the operand is directly addressed.</p> <p>Format = U3</p> <p>0x00 to 0x07 = Address Sub-Register number in the range of [0, 7]</p>	<p>3</p>



Field	Description	Width
<b>AddrImm</b>	<p><b>Address Immediate.</b> This field provides the immediate value in unit of byte to be added to the address register in order to compute the register address (byte-aligned region origin) for the operand. It is a 10-bit signed integer in 2's compliment form. This field is present if the operand is in register-indirect addressing mode; it is not present if the operand is directly addressed.</p> <p><i>Note: that the address immediate field may not be able to cover the whole GRF register range for a thread, as the maximum GRF register space for a thread is 4KB.</i></p> <p>Format = S9 Valid range: [-512, 511]</p>	10
<b>SrcMod</b>	<p><b>Source Modifier.</b> This field specifies the numerical modification to a source operand. The value of each data element of a source operand can optionally have its absolute value taken and/or its sign inverted prior to delivery to the execution pipe. The absolute value is prior to negate such that a guaranteed negative value can be produced.</p> <p>This field only applies to source operand. It does not apply to destination. This field is not present for an immediate source operand.</p> <p>00 = No modification (normal) 01 = “<b>(abs)</b>”. Absolute 10 = “-”. Negate 11 = “-<b>(abs)</b>”. Negate of the absolute (forced negative value)</p>	2



Field	Description	Width
<p><b>VertStride</b></p>	<p><b>Vertical Stride.</b> The field provides the vertical stride of the register region in unit of data elements for an operand.</p> <p>Encoding of this field provides values in power of 2, ranging from 0 to 32 elements. Larger values are not supported due to the restriction that a source operand must reside within two adjacent 256-bit registers (64 bytes total).</p> <p>Special encoding 1111b (0xF) is only valid when the operand is in register-indirect addressing mode (<i>AddrMode</i> = 1). If this field is set to 0xF, one or more sub-registers of the address registers may be used to compute the addresses. Each address sub-register provides the origin for a row of data element. The number of address sub-registers used is determined by the division of <i>ExecSize</i> of the instruction by the <i>Width</i> fields of the operand.</p> <p>This field only applies to source operand. It does not apply to destination.</p> <p>This field is not present for an immediate source operand.</p> <p>For <b>Align16</b> access mode, only encodings of 0000 and 0011 are allowed. Other codes are reserved.</p> <p><i>Note 1: Vertical Stride larger than 32 is not allowed due to the restriction that a source operand must reside within two adjacent 256-bit registers (64 bytes total).</i></p> <p><i>Note 2: In <b>Align16</b> access mode, as encoding 0xF is reserved, only single-index indirect addressing is supported.</i></p> <p><i>Note 3: If indirect address is supported for &lt;src1&gt;, encoding 0xF is reserved for &lt;src1&gt; – only single-index indirect addressing is supported.</i></p> <p>0000 = 0 Elements  0001 = 1 Element  0010 = 2 Elements  0011 = 4 Elements  0100 = 8 Elements  0101 = 16 Elements (applies to byte or word operand only)  0110 = 32 Elements (applies to byte operand only)  0111-1110 = Reserved  1111 = <b>VxH or Vx1 mode</b> (only valid for register-indirect addressing in <b>Align1</b> mode)</p>	<p>4</p>
<p><b>Width</b></p>	<p><b>Width.</b> This field specifies the number of elements in the horizontal dimension of the region for a source operand. This field cannot exceed the <i>ExecSize</i> field of the instruction.</p> <p>This field only applies to source operand. It does not apply to destination.</p> <p>This field is not present for an immediate source operand.</p> <p>000 = 1 Elements  001 = 2 Elements  010 = 4 Elements  011 = 8 Elements  100 = 16 Elements  101-111 = Reserved</p>	<p>3</p>



Field	Description	Width																
<b>HorzStride</b>	<p><b>Horizontal Stride.</b> This field provides the distance in unit of data elements between two adjacent data elements within a row (horizontal) in the register region for the operand.</p> <p>This field applies to both destination and source operands.</p> <p>This field is not present for an immediate source operand.</p> <p>00 = 0 Elements 01 = 1 Element 10 = 2 Elements 11 = 4 Elements</p>	2																
<b>Imm32</b>	<p><b>32-bit Immediate.</b> The 32-bit immediate data field for the operand. It may contain any legal bit pattern for its associated type. Only one 32-bit immediate value may be present in an instruction, therefore binary operations only support &lt;src1&gt; as an immediate value.</p> <p>The low order bits are directly used when fewer than 32-bits are needed to describe the desired type; the 32-bits are not coerced into the designated type.</p> <p>For UW and W data types, programmer is required to replicate the lower word to the upper word of this field.</p> <p>This field only applies to the last source operand.</p> <p>Signed and unsigned byte integer data types are not supported for an immediate operand.</p> <p>Valid ranges according to data type:</p> <table border="1" data-bbox="485 1020 1289 1388"> <thead> <tr> <th data-bbox="485 1020 777 1073">Immediate Data Type</th> <th data-bbox="777 1020 1289 1073">Valid Range (inclusive)</th> </tr> </thead> <tbody> <tr> <td data-bbox="485 1073 777 1125">F</td> <td data-bbox="777 1073 1289 1125">[0...±1.0*2<sup>-128...127</sup>]</td> </tr> <tr> <td data-bbox="485 1125 777 1167">UW</td> <td data-bbox="777 1125 1289 1167">[0, 65535]</td> </tr> <tr> <td data-bbox="485 1167 777 1209">W</td> <td data-bbox="777 1167 1289 1209">[-32768, 32767]</td> </tr> <tr> <td data-bbox="485 1209 777 1251">UD</td> <td data-bbox="777 1209 1289 1251">[0, 2<sup>32</sup>-1]</td> </tr> <tr> <td data-bbox="485 1251 777 1293">D</td> <td data-bbox="777 1251 1289 1293">[-2<sup>31</sup>, 2<sup>31</sup>-1]</td> </tr> <tr> <td data-bbox="485 1293 777 1335">VF</td> <td data-bbox="777 1293 1289 1335">[0, ±0.125...±31]</td> </tr> <tr> <td data-bbox="485 1335 777 1388">V</td> <td data-bbox="777 1335 1289 1388">[-8, 7]</td> </tr> </tbody> </table>	Immediate Data Type	Valid Range (inclusive)	F	[0...±1.0*2 <sup>-128...127</sup> ]	UW	[0, 65535]	W	[-32768, 32767]	UD	[0, 2 <sup>32</sup> -1]	D	[-2 <sup>31</sup> , 2 <sup>31</sup> -1]	VF	[0, ±0.125...±31]	V	[-8, 7]	32
Immediate Data Type	Valid Range (inclusive)																	
F	[0...±1.0*2 <sup>-128...127</sup> ]																	
UW	[0, 65535]																	
W	[-32768, 32767]																	
UD	[0, 2 <sup>32</sup> -1]																	
D	[-2 <sup>31</sup> , 2 <sup>31</sup> -1]																	
VF	[0, ±0.125...±31]																	
V	[-8, 7]																	
<b>ChanEn</b>	<p><b>Channel Enable.</b> Four channel enables are defined for controlling which channels will be written into the destination region. These channel mask bits are applied in a modulo-four manner to all <i>ExecSize</i> channels. There is 1-bit Channel Enable for each channel within the group of 4. If the bit is cleared, the write for the corresponding channel is disabled. If the bit is set, the write is enabled. Mnemonic for the bit being set for the group of 4 is “x”, “y”, “z”, and “w”, respectively, where “x” corresponds to Channel 0 in the group and “w” corresponds to channel 3 in the group.</p> <p>This field only applies to destination operand.</p> <p>This field is only present in <b>Align16</b> mode.</p> <p>0 = Write Disabled 1 = Write Enabled (normal)</p>	4																





Field	Description	Width
<b>ChanSel</b>	<p><b>Channel Select.</b> This field controls the channel swizzle for a source operand. The normally sequential channel assignment can be altered by explicitly identifying neighboring data elements for each channel. Out of the 8-bit field, 2 bits are assigned for each channel within the group of 4. ChanSel[1:0], [3.2], [5.4] and [7.6] are for channel 0 (“x”), 1 (“y”), 2 (“z”), and 3 (“w”) in the group, respectively.</p> <p>For example with an execution size of 8, <i>r0.0&lt;4&gt;.zywz:f</i> would assign the channels as follows: Chan<sub>0</sub> = Data<sub>2</sub>, Chan<sub>1</sub> = Data<sub>1</sub>, Chan<sub>2</sub> = Data<sub>3</sub>, Chan<sub>3</sub> = Data<sub>2</sub>; Chan<sub>4</sub> = Data<sub>6</sub>, Chan<sub>5</sub> = Data<sub>5</sub>, Chan<sub>6</sub> = Data<sub>7</sub>, Chan<sub>7</sub> = Data<sub>6</sub>.</p> <p>This field only applies to source operand.</p> <p>This field is only present in <b>Align16</b> mode. It is not present for an immediate source operand.</p> <p>The 2-bit Channel Selection field for each channel within the group of 4 is defined as the following.</p> <ul style="list-style-type: none"> <li>00 = “x”. Channel 0 is selected for the corresponding execution channel</li> <li>01 = “y”. Channel 1 is selected for the corresponding execution channel</li> <li>10 = “z”. Channel 2 is selected for the corresponding execution channel</li> <li>11 = “w”. Channel 3 is selected for the corresponding execution channel</li> </ul>	8
<b>MsgDscpt31</b>	<p><b>Message Description.</b> This field, containing 31-bit immediate values, provides the description of the message to be sent.</p> <p>This field only applies to the <i>send</i> instruction. It is not present for other instructions.</p> <p>The meaning of the field depends on the type of message as well as the message shared function target.</p> <p>Format: U31</p>	31
<b>EOT</b>	<p><b>End of Thread.</b> This field controls the termination of the thread. For a <i>send</i> instruction, if this field is set, EU will terminate the thread and also set the EOT bit in the message sideband.</p> <p>This field only applies to the <i>send</i> instruction. It is not present for other instructions.</p> <ul style="list-style-type: none"> <li>0 = The thread is not terminated</li> <li>1 = <b>EOT</b></li> </ul>	1



### 7.2.1.2 3-src Instructions

The table in this section describes the encoding for the common fields for the 3-src instructions format.

**Table 7-3. Definitions of Common Instruction Fields**

Field	Description	Width
<b>CondModifier</b>	<p><b>Conditional Modifier.</b> This field sets the flag register based on the internal conditional signals output from the execution pipe such as sign, zero, overflow and NaNs, etc. If this field is set to 0000, no flag registers are updated. Flag registers are not updated for instructions with embedded compares.</p> <p>This field may also be referred to as the <i>flag destination control</i> field.</p> <p>This field applies to all instructions except <i>send</i>, <i>sendc</i>, and <i>math</i>.</p> <p>0000 = Do not modify the flag register (normal)            0001 = Zero or Equal (‘.z’ or ‘.e’)            0010 = Not Zero or Not Equal (‘.nz’ or ‘.ne’)            0011 = Greater-than (‘.g’)            0100 = Greater-than-or-equal (‘.ge’)            0101 = Less-than (‘.l’)            0110 = Less-than-or-equal (‘.le’)            0111 = Reserved            1000 = Overflow (‘.o’)            1001 = Unordered with Computed NaN (‘.u’)            1010 -1111 = Reserved</p>	4
<b>RegNum</b>	<p><b>Register Number.</b> This field provides the register number for the operand. For GRF or MRF register operand, it provides the portion of register address aligning to 256-bit. For an ARF register operand, this field is encoded such that MSBs identify the architecture register type and LSBs provide its register number.</p> <p>This field together with the corresponding <i>SubRegNum</i> field provides the byte aligned address for the origin of the register region. Specifically, this field provides bits [12:5] of the byte address, while <i>SubRegNum</i> field provides bits [4:0].</p> <p>This field applies to the destination operand and the source operands. It is ignored (or not present in the instruction word) for an immediate source operand.</p> <p>This field is present if the operand is in direct addressing mode; it is not present if the operand is register-indirect addressed.</p> <p>Format = U8, if <i>RegFile</i> = <b>GRF</b>.</p> <p>0x00 to 0x7F = Register number in the range of [0, 127]            0x80 to 0xFF = Reserved</p> <p>Format = U8, if <i>RegFile</i> = <b>MRF</b>.</p> <p>0x00 to 0x0F = Register number in the range of [0, 15]            0x10 to 0xFF = Reserved</p>	8
<b>SubRegNum</b>	<p><b>Sub-Register Number.</b> This field provides the sub-register number for the operand, it provides the <b>dword</b> address within a 256-bit register</p> <p>This field together with the corresponding <i>RegNum</i> field provides the dword aligned address for the origin of the register region. Specifically, this field provides bits [4:2] of the dword address, while <i>RegNum</i> field provides bits [12:5].</p> <p>This field applies to the destination operand and the source operands.</p>	3



Field	Description	Width
<b>SrcMod</b>	<p><b>Source Modifier.</b> This field specifies the numerical modification to a source operand. The value of each data element of a source operand can optionally have its absolute value taken and/or its sign inverted prior to delivery to the execution pipe. The absolute value is prior to negate such that a guaranteed negative value can be produced.</p> <p>This field only applies to source operand. It does not apply to destination.</p> <p>This field is not present for an immediate source operand.</p> <p>00 = No modification (normal)            01 = “<b>(abs)</b>”. Absolute            10 = “<b>-</b>”. Negate            11 = “<b>-(abs)</b>”. Negate of the absolute (forced negative value)</p>	2
<b>ChanSel</b>	<p><b>Channel Select.</b> This field controls the channel swizzle for a source operand. The normally sequential channel assignment can be altered by explicitly identifying neighboring data elements for each channel. Out of the 8-bit field, 2 bits are assigned for each channel within the group of 4. ChanSel[1:0], [3,2], [5,4] and [7,6] are for channel 0 (“x”), 1 (“y”), 2 (“z”), and 3 (“w”) in the group, respectively.</p> <p>For example with an execution size of 8, <i>r0.0&lt;4&gt;.zywz:f</i> would assign the channels as follows: Chan<sub>0</sub> = Data<sub>2</sub>, Chan<sub>1</sub> = Data<sub>1</sub>, Chan<sub>2</sub> = Data<sub>3</sub>, Chan<sub>3</sub> = Data<sub>2</sub>; Chan<sub>4</sub> = Data<sub>6</sub>, Chan<sub>5</sub> = Data<sub>5</sub>, Chan<sub>6</sub> = Data<sub>7</sub>, Chan<sub>7</sub> = Data<sub>6</sub>.</p> <p>This field only applies to source operand.</p> <p>This field is only present in <b>Align16</b> mode. It is not present for an immediate source operand.</p> <p>The 2-bit Channel Selection field for each channel within the group of 4 is defined as the following.</p> <p>00 = “<b>x</b>”. Channel 0 is selected for the corresponding execution channel            01 = “<b>y</b>”. Channel 1 is selected for the corresponding execution channel            10 = “<b>z</b>”. Channel 2 is selected for the corresponding execution channel            11 = “<b>w</b>”. Channel 3 is selected for the corresponding execution channel</p>	8
<b>RepCtrl</b>	<p><b>Replicate Control.</b> This field controls the replication of the starting channel to all channels in the execution size.</p> <p>This field applies to all three source operands.</p> <p>0 = No replication            1 = Replicate across all channels</p>	1

## 7.2.2 Instruction Operation Doubleword (DW0)

Most fields in Instruction Operation Doubleword (DW0) apply to all instructions. Bit field [27:24] is one exception. It is *CondModifier* for most instructions but is *CurrDest.RegNum* field for the *send* instruction.

The descriptions in the table below are shared between the 1-src/2-src instructions and 3-src instructions.



**Table 7-4. Definitions of Fields in Operation Doubleword (DW0)**

Bits	Description																
31	<p><b>Saturate.</b> This field controls the destination saturation.</p> <p>When it is set, output data to the destination register are saturated. The saturation operation depends on the destination data type. Saturation is the operation that converts any data that is outside the <i>saturation target range</i> for the data type to the closest represented value with the target range. If destination type is float, saturation target range is [0, 1]. For example, any positive number greater than 1 (including +INF) is saturated to 1 and any negative number (including -INF) is saturated to 0. A NaN is saturated to 0. For integer data types, the maximum range for the given numerical data type is the saturation target range.</p> <p>When it is not set, output data to the destination register are not saturated. For example, a wrapped result (modular) is output to the destination for an overflowed integer data.</p> <p>More details can be found in the Data Types chapter.</p> <p>0 = No destination modification (normal)            1 = “<b>sat</b>”. Saturate the output</p> <table border="1" style="margin-left: 40px; margin-top: 10px;"> <thead> <tr> <th style="text-align: center;">Destination Type</th> <th style="text-align: center;">Saturation Target Range (inclusive)</th> </tr> </thead> <tbody> <tr> <td>Float (F)</td> <td>[0.0, 1.0]</td> </tr> <tr> <td>Byte (UB)</td> <td>[0, 255]</td> </tr> <tr> <td>Signed Byte (B)</td> <td>[-128, 127]</td> </tr> <tr> <td>Word (UW)</td> <td>[0, 65535]</td> </tr> <tr> <td>Signed Word (W)</td> <td>[-32768, 32767]</td> </tr> <tr> <td>Double Word (UD)</td> <td>[0, 2<sup>32</sup>-1]</td> </tr> <tr> <td>Signed Double (D)</td> <td>[-2<sup>31</sup>, 2<sup>31</sup>-1]</td> </tr> </tbody> </table>	Destination Type	Saturation Target Range (inclusive)	Float (F)	[0.0, 1.0]	Byte (UB)	[0, 255]	Signed Byte (B)	[-128, 127]	Word (UW)	[0, 65535]	Signed Word (W)	[-32768, 32767]	Double Word (UD)	[0, 2 <sup>32</sup> -1]	Signed Double (D)	[-2 <sup>31</sup> , 2 <sup>31</sup> -1]
Destination Type	Saturation Target Range (inclusive)																
Float (F)	[0.0, 1.0]																
Byte (UB)	[0, 255]																
Signed Byte (B)	[-128, 127]																
Word (UW)	[0, 65535]																
Signed Word (W)	[-32768, 32767]																
Double Word (UD)	[0, 2 <sup>32</sup> -1]																
Signed Double (D)	[-2 <sup>31</sup> , 2 <sup>31</sup> -1]																
30	<b>Reserved</b>																
29	Reserved: MBZ																
28	<p><b>AccWrCtrl.</b> This field allows per instruction accumulator write control.</p> <p>0 = don't write result into accumulator            1 = “<b>AccWrCtrl</b>”. write result into accumulator, and destination</p>																
27:24	<p><b>CondModifier or CurrDst.RegNum[3:0]</b></p> <p>Definition of this bit field depends on whether the instruction is a <i>send/math</i> or not.</p> <table border="1" style="margin-left: 40px; margin-top: 10px;"> <thead> <tr> <th style="text-align: center;">Opcode != 'send'</th> <th style="text-align: center;">Opcode = 'send'</th> </tr> </thead> <tbody> <tr> <td> <p><b>CondModifier:</b></p> <p>This field sets the flag register based on the internal conditional signals output from the execution pipe.</p> </td> <td> <p><b>CurrDst.RegNum[3:0]</b></p> <p>This field sets the MRF register number for the current destination operand in the <i>send</i> instruction. No flag registers are updated for the <i>send</i> instruction. The 4-bit field provides full access of the 16 MRF registers.            (See Instruction Reference chapter for <i>CurrDst</i>.)</p> </td> </tr> </tbody> </table>	Opcode != 'send'	Opcode = 'send'	<p><b>CondModifier:</b></p> <p>This field sets the flag register based on the internal conditional signals output from the execution pipe.</p>	<p><b>CurrDst.RegNum[3:0]</b></p> <p>This field sets the MRF register number for the current destination operand in the <i>send</i> instruction. No flag registers are updated for the <i>send</i> instruction. The 4-bit field provides full access of the 16 MRF registers.            (See Instruction Reference chapter for <i>CurrDst</i>.)</p>												
Opcode != 'send'	Opcode = 'send'																
<p><b>CondModifier:</b></p> <p>This field sets the flag register based on the internal conditional signals output from the execution pipe.</p>	<p><b>CurrDst.RegNum[3:0]</b></p> <p>This field sets the MRF register number for the current destination operand in the <i>send</i> instruction. No flag registers are updated for the <i>send</i> instruction. The 4-bit field provides full access of the 16 MRF registers.            (See Instruction Reference chapter for <i>CurrDst</i>.)</p>																



Bits	Description
23:21	<p><b>ExecSize – Execution Size.</b> This field determines the number of channels operating in parallel for this instruction. The size cannot exceed the maximum number of channels allowed for the given data type.</p> <ul style="list-style-type: none"> <li>000 = 1 Channels (scalar operation)</li> <li>001 = 2 Channels</li> <li>010 = 4 Channels</li> <li>011 = 8 Channels</li> <li>100 = 16 Channels</li> <li>101 = 32 Channels</li> <li>110-111 = Reserved</li> </ul>
20	<p><b>PredInv – Predicate Inverse.</b> This field, together with <i>PredCtrl</i>, enables and controls the generation of the predication mask for the instruction. When it is set, the predication uses the inverse of the predication bits generated according to setting of Predicate Control. In other words, effect of PredInv happens after PredCtrl.</p> <p>This field is ignored by hardware if Predicate Control is set to 0000 – there is no predication.</p> <ul style="list-style-type: none"> <li>0 = “+”. Positive polarity of predication.</li> <li>1 = “-”. Negative polarity of predication.</li> </ul>
19:16	<p><b>PredCtrl – Predicate Control.</b> This field, together with <i>PredInv</i>, enables and controls the generation of the predication mask for the instruction. It allows per-channel conditional execution of the instruction based on the content of the selected flag register. Encoding depends on the access mode.</p> <p>In <b>Align16</b> access mode, there are eight encodings (including no predication). All encodings are based on group-of-4 predicate bits, including channel sequential, replication swizzles and horizontal any/all operations. The same configuration is repeated for each group-of-4 execution channels.</p> <p>In <b>Align1</b> access mode, there are twelve encodings (including no predication). The encodings applies to all execution channels with explicit channel grouping from single channel up to group of 16 channels.</p> <p>Predicate Control in <b>Align16</b> access mode</p> <ul style="list-style-type: none"> <li>0000 = No predication (normal)</li> <li>0001 = Predication with sequential flag channel mapping</li> <li>0010 = Predication with replication swizzle ‘.x’</li> <li>0011 = Predication with replication swizzle ‘.y’</li> <li>0100 = Predication with replication swizzle ‘.z’</li> <li>0101 = Predication with replication swizzle ‘.w’</li> <li>0110 = Predication with ‘.any4h’</li> <li>0111 = Predication with ‘.all4h’</li> <li>1000 -1111 = Reserved</li> </ul> <p>Predicate Control in <b>Align1</b> access mode</p> <ul style="list-style-type: none"> <li>0000 = No predication (normal)</li> <li>0001 = Predication with sequential flag channel mapping</li> <li>0010 = Predication with <b>.anyv</b> (any from f0.0-f0.1 on the same channel)</li> <li>0011 = Predication with <b>.allv</b> (all of f0.0-f0.1 on the same channel)</li> <li>0100 = Predication with <b>.any2h</b> (any in group of 2 channels)</li> <li>0101 = Predication with <b>.all2h</b> (all in group of 2 channels)</li> <li>0110 = Predication with <b>.any4h</b> (any in group of 4 channels)</li> <li>0111 = Predication with <b>.all4h</b> (all in group of 4 channels)</li> <li>1000 = Predication with <b>.any8h</b> (any in group of 8 channels)</li> <li>1001 = Predication with <b>.all8h</b> (all in group of 8 channels)</li> </ul>



Bits	Description																																																												
	1010 = Predication with <b>.any16h</b> (any in group of 16 channels) 1011 = Predication with <b>.all16h</b> (all in group of 16 channels) 1100 = Reserved 1101 = Reserved 1100 -1111 = Reserved [ <b>DevSNB</b> ]																																																												
15:14	<p><b>ThreadCtrl – Thread Control.</b> This field provides explicit control for thread switching.</p> <p>If this field is set to 00, it is up to the execution units to manage thread switching. This is the normal operations mode. In this mode, for example, if the current instruction cannot proceed due to operand dependencies, EU switches to next available thread to fill the compute pipe. In another example, if the current instruction is ready to go, however, there is another thread with higher priority also has instruction ready, EU switches to that thread.</p> <p>If this field is set to <b>Switch</b>, a forced thread switch occurs after the current instruction is executed and before the next instruction. In addition, a long delay (longer than the execution pipe latency) for the current thread is introduced for the thread. Particularly, the instruction queue of the current thread is flushed after the current instruction is dispatched for execution.</p> <p>If this field is set to <b>Atomic</b>, the next instruction will get highest priority in the thread arbitration for the execution pipelines.</p> <p><b>Switch</b> is designed primarily as a safety feature in case there are race conditions for certain instructions.</p> <ul style="list-style-type: none"> <li>00 = Normal Thread Control</li> <li>10 = “<b>Switch</b>”</li> <li>01 = “<b>Atomic</b>”</li> <li>11 = Reserved</li> </ul>																																																												
13:12	<p><b>QtrCtrl – Quarter Control.</b> This field provides explicit control for ARF selection.</p> <p>This field combines with <b>ExecSize</b> determines which channels are used for the ARF registers.</p> <table border="1" data-bbox="337 1129 1360 1875"> <thead> <tr> <th data-bbox="337 1129 464 1171">QtrCtrl</th> <th data-bbox="464 1129 607 1171">NibCtrl</th> <th data-bbox="607 1129 750 1171">ExecSize</th> <th data-bbox="750 1129 1253 1171">Description</th> <th data-bbox="1253 1129 1360 1171">BNF</th> </tr> </thead> <tbody> <tr> <td data-bbox="337 1171 464 1255">00</td> <td data-bbox="464 1171 607 1255">x</td> <td data-bbox="607 1171 750 1255">8</td> <td data-bbox="750 1171 1253 1255">use first quarter for DMask/VMask use first half for everything else</td> <td data-bbox="1253 1171 1360 1255"><b>1Q</b></td> </tr> <tr> <td data-bbox="337 1255 464 1339">01</td> <td data-bbox="464 1255 607 1339">x</td> <td data-bbox="607 1255 750 1339">8</td> <td data-bbox="750 1255 1253 1339">use second quarter for DMask/VMask use second half for everything else</td> <td data-bbox="1253 1255 1360 1339"><b>2Q</b></td> </tr> <tr> <td data-bbox="337 1339 464 1423">10</td> <td data-bbox="464 1339 607 1423">x</td> <td data-bbox="607 1339 750 1423">8</td> <td data-bbox="750 1339 1253 1423">use third quarter for DMask/VMask use first half for everything else</td> <td data-bbox="1253 1339 1360 1423"><b>3Q</b></td> </tr> <tr> <td data-bbox="337 1423 464 1507">11</td> <td data-bbox="464 1423 607 1507">x</td> <td data-bbox="607 1423 750 1507">8</td> <td data-bbox="750 1423 1253 1507">use forth quarter for DMask/VMask use second half for everything else</td> <td data-bbox="1253 1423 1360 1507"><b>4Q</b></td> </tr> <tr> <td data-bbox="337 1507 464 1591">0x</td> <td data-bbox="464 1507 607 1591">x</td> <td data-bbox="607 1507 750 1591">16</td> <td data-bbox="750 1507 1253 1591">use first half for DMask/VMask use all channels for everything else</td> <td data-bbox="1253 1507 1360 1591"><b>1H</b></td> </tr> <tr> <td data-bbox="337 1591 464 1675">1x</td> <td data-bbox="464 1591 607 1675">x</td> <td data-bbox="607 1591 750 1675">16</td> <td data-bbox="750 1591 1253 1675">use second half for DMask/VMask use all channels for everything else</td> <td data-bbox="1253 1591 1360 1675"><b>2H</b></td> </tr> <tr> <td data-bbox="337 1675 464 1707">00</td> <td data-bbox="464 1675 607 1707">0</td> <td data-bbox="607 1675 750 1707">4</td> <td data-bbox="750 1675 1253 1707">use first 1/8 for DMask/VMask and ARF</td> <td data-bbox="1253 1675 1360 1707"><b>1N</b></td> </tr> <tr> <td data-bbox="337 1707 464 1738">00</td> <td data-bbox="464 1707 607 1738">1</td> <td data-bbox="607 1707 750 1738">4</td> <td data-bbox="750 1707 1253 1738">use second 1/8 for DMask/VMask and ARF</td> <td data-bbox="1253 1707 1360 1738"><b>2N</b></td> </tr> <tr> <td data-bbox="337 1738 464 1770">01</td> <td data-bbox="464 1738 607 1770">0</td> <td data-bbox="607 1738 750 1770">4</td> <td data-bbox="750 1738 1253 1770">use thrid 1/8 for DMask/VMask and ARF</td> <td data-bbox="1253 1738 1360 1770"><b>3N</b></td> </tr> <tr> <td data-bbox="337 1770 464 1801">01</td> <td data-bbox="464 1770 607 1801">1</td> <td data-bbox="607 1770 750 1801">4</td> <td data-bbox="750 1770 1253 1801">use forth 1/8 for DMask/VMask and ARF</td> <td data-bbox="1253 1770 1360 1801"><b>4N</b></td> </tr> <tr> <td data-bbox="337 1801 464 1833">10</td> <td data-bbox="464 1801 607 1833">0</td> <td data-bbox="607 1801 750 1833">4</td> <td data-bbox="750 1801 1253 1833">use fifth 1/8 for DMask/VMask and ARF</td> <td data-bbox="1253 1801 1360 1833"><b>5N</b></td> </tr> </tbody> </table>	QtrCtrl	NibCtrl	ExecSize	Description	BNF	00	x	8	use first quarter for DMask/VMask use first half for everything else	<b>1Q</b>	01	x	8	use second quarter for DMask/VMask use second half for everything else	<b>2Q</b>	10	x	8	use third quarter for DMask/VMask use first half for everything else	<b>3Q</b>	11	x	8	use forth quarter for DMask/VMask use second half for everything else	<b>4Q</b>	0x	x	16	use first half for DMask/VMask use all channels for everything else	<b>1H</b>	1x	x	16	use second half for DMask/VMask use all channels for everything else	<b>2H</b>	00	0	4	use first 1/8 for DMask/VMask and ARF	<b>1N</b>	00	1	4	use second 1/8 for DMask/VMask and ARF	<b>2N</b>	01	0	4	use thrid 1/8 for DMask/VMask and ARF	<b>3N</b>	01	1	4	use forth 1/8 for DMask/VMask and ARF	<b>4N</b>	10	0	4	use fifth 1/8 for DMask/VMask and ARF	<b>5N</b>
QtrCtrl	NibCtrl	ExecSize	Description	BNF																																																									
00	x	8	use first quarter for DMask/VMask use first half for everything else	<b>1Q</b>																																																									
01	x	8	use second quarter for DMask/VMask use second half for everything else	<b>2Q</b>																																																									
10	x	8	use third quarter for DMask/VMask use first half for everything else	<b>3Q</b>																																																									
11	x	8	use forth quarter for DMask/VMask use second half for everything else	<b>4Q</b>																																																									
0x	x	16	use first half for DMask/VMask use all channels for everything else	<b>1H</b>																																																									
1x	x	16	use second half for DMask/VMask use all channels for everything else	<b>2H</b>																																																									
00	0	4	use first 1/8 for DMask/VMask and ARF	<b>1N</b>																																																									
00	1	4	use second 1/8 for DMask/VMask and ARF	<b>2N</b>																																																									
01	0	4	use thrid 1/8 for DMask/VMask and ARF	<b>3N</b>																																																									
01	1	4	use forth 1/8 for DMask/VMask and ARF	<b>4N</b>																																																									
10	0	4	use fifth 1/8 for DMask/VMask and ARF	<b>5N</b>																																																									



Bits	Description				
	10	1	4	use sixth 1/8 for DMask/VMask and ARF	<b>6N</b>
	11	0	4	use seventh 1/8 for DMask/VMask and ARF	<b>7N</b>
	11	1	4	use eighth 1/8 for DMask/VMask and ARF	<b>8N</b>
11:10	<p><b>DepCtrl – Destination Dependency Control.</b> This field selectively disables destination dependency check and clear for this instruction.</p> <p>When it is set to 00, normal destination dependency control is performed for the instruction – hardware checks for destination hazards to ensure data integrity. Specifically, destination register dependency check is conducted before the instruction is made ready for execution. After the instruction is executed, the destination register scoreboard will be cleared when the destination operands retire.</p> <p>When bit 10 is set (<b>NoDDClr</b>), the destination register scoreboard will NOT be cleared when the destination operands retire. When bit 11 is set (<b>NoDDChk</b>), hardware does not check for destination register dependency before the instruction is made ready for execution. <b>NoDDClr</b> and <b>NoDDChk</b> are not mutual exclusive.</p> <p>When this field is not all-zero, hardware does not protect against destination hazards for the instruction. This is typically used to assemble data in a fine grained fashion (e.g. matrix-vector compute with dot-product instructions), where the data integrity is guaranteed by software based on the intended usage of instruction sequences.</p> <p>00 = Destination dependency checked and cleared (normal)  01 = “<b>NoDDClr</b>”. Destination dependency checked but not cleared  10 = “<b>NoDDChk</b>”. Destination dependency not checked but cleared  11 = “<b>NoDDClr, NoDDChk</b>”. Destination dependency not checked and not cleared</p> <p>[Errata DevSNB] DepCtrl is not supported for ARF registers.</p>				
9	<p><b>WECtrl – Write Enable Control.</b> This field determines if the the per channel write enables are used to generate the final write enable. This field should be normally “0”.</p> <p>0 = use normal write enables (normal)  1 = write all channels, except channels killed with predication control</p>				
8	<p><b>AccessMode – Access Mode.</b> This field determines the operand access for the instruction. It applies to all source and destination operands.</p> <p>When it is cleared (<b>Align1</b>), the instruction uses byte-aligned addressing for source and destination operands. Source swizzle control and destination mask control are not supported.</p> <p>When it is set (<b>Align16</b>), the instruction uses 16-byte-aligned addressing for all source and destination operands. Source swizzle control and destination mask control are supported in this mode.</p> <p>0 = “<b>Align1</b>”  1 = “<b>Align16</b>”</p>				
7	Reserved: MBZ (for future opcode extension)				
6:0	<p><b>Opcode – Instruction Operation Code.</b> This field contains the instruction operation code. Each opcode is given a unique mnemonic. For example, opcode 0x01 is for a move operation. Mnemonic for this opcode is <i>mov</i>.</p> <p>See Section 7.3 for details of opcode encoding.</p>				



## 7.2.3 Instruction Destination Doubleword (DW1)

### 7.2.3.1 1-src and 2-src Instructions

Destination Doubleword (DW1) contains the register file and numeric type of all operands, as well as the register region parameters of the destination operand. Table 7-5 shows the field definition of the Instruction Destination Doubleword. Furthermore, the Destination Register Region is described in Table 7-6 through Table 7-9.

**Table 7-5. Instruction Destination Doubleword**

Bits	Description
31:16	<p><b>Destination Register Region.</b> This word contains the parameters describing the register region of the destination operand. Subfield definition depends on the AccessMode. Detailed descriptions can be found in Table 7-6 through Table 7-9.</p> <p><b>Programming Notes:</b> Although <i>Dst.HorzStride</i> is a don't care for Align16, HW needs this to be programmed as "01".</p>
15	Reserved: MBZ
14:12	<p><b>Src1.SrcType – Source-1 Data Type.</b> This field specifies the numerical data type of the source operand &lt;src1&gt;. The bits of a source operand are interpreted as the identified numerical data type, rather than coerced into a type implied by the operator. Depending on <i>RegFile</i> field of the source operand, there are two different encoding for this field. If a source is a register operand, this field follows the Source Register Type Encoding. If a source is an immediate operand, this field follows the Source Immediate Type Encoding.</p> <p>Source Register Type Encoding is identical to that for Destination Type.</p> <p>Source Immediate Type Encoding differs in two areas. First, it does not support byte and unsigned numerical data types. Secondly, it has two 32-bit vector types – halfbyte integer vector (V) type and exponent-only float vector (VF) type.</p> <p><i>Implementation Note 1: Both source operands, &lt;src0&gt; and &lt;src1&gt;, support immediate types, but only one immediate is allowed for a given instruction and it must be the last operand.</i></p> <p><i>Implementation Note 2: Halfbyte integer vector (v) type can only be used in instructions in packed-word execution mode. Therefore, in a two-source instruction where &lt;src1&gt; is of type :v, &lt;src0&gt; must be of type :b, :ub, :w, or :uw.</i></p> <p>Source Register Type Encoding</p> <ul style="list-style-type: none"> <li>000 = “UD”. Unsigned Doubleword integer</li> <li>001 = “D”. Signed Doubleword integer</li> <li>010 = “UW”. Unsigned Word integer</li> <li>011 = “W”. Signed Word integer</li> <li>100 = “UB”. Unsigned Byte integer</li> <li>101 = “B”. Signed Byte integer</li> <li>110 = Reserved [DevSNB]</li> <li>111 = “F”. Single precision Float (32-bit)</li> </ul> <p>Source Immediate Type Encoding:</p> <ul style="list-style-type: none"> <li>000 = “UD”</li> <li>001 = “D”</li> <li>010 = “UW”</li> <li>011 = “W”</li> </ul>





Bits	Description
	100 = “UV”. 32-bit halfbyte <b>Unsigned Integer Vector</b> 101 = “VF”. 32-bit restricted <b>Vector Float</b> 110 = “V”. 32-bit halfbyte integer <b>Vector</b> 111 = “F”
11:10	<b>Src1.RegFile – Source-1 Register File.</b> This field identifies the register file of source operand <src1>. <ul style="list-style-type: none"> <li>00 = “ARF”. Architecture Register File (<b>a#</b>, <b>acc#</b>, <b>f#</b>, <b>n#</b>, <b>null</b>, <b>ip</b>, etc.)</li> <li>01 = “GRF”. General Register File (<b>r#</b>)</li> <li>10 = “MRF”. Message Register File (<b>m#</b>)</li> <li>11 = “IMM”. Immediate</li> </ul>
9:7	<b>Src0.SrcType – Source-0 Data Type.</b> This field is the <i>SrcType</i> for <src0> operand. It has the same definitions as <i>Src1.SrcType</i> .
6:5	<b>Src0.RegFile – Source-0 Register File.</b> This field is the <i>RegFile</i> for <src0> operand. It has the same definitions as <i>Src1.RegFile</i> .
4:2	<b>Dst.DstType – Destination Data Type.</b> This field specifies the numerical data type of the destination operand <dst>. The bits of the destination operand are interpreted as the identified numerical data type, rather than coerced into a type implied by the operator. For a <i>send</i> instruction, this field applies to the <b>CurrDst</b> – the current destination operand. Encoding: <ul style="list-style-type: none"> <li>000 = “UD”. <b>Unsigned Doubleword integer</b></li> <li>001 = “D”. <b>Signed Doubleword integer</b></li> <li>010 = “UW”. <b>Unsigned Word integer</b></li> <li>011 = “W”. <b>Signed Word integer</b></li> <li>100 = “UB”. <b>Unsigned Byte integer</b></li> <li>101 = “B”. <b>Signed Byte integer</b></li> <li>110 = Reserved</li> <li>111 = “F”. <b>Single precision Float (32-bit)</b></li> </ul>
1:0	<b>Dst.RegFile – Destination Register File.</b> This field identifies the register file of the destination operand <dst>. Note that it is obvious that immediate cannot be a destination operand. For a <i>send</i> instruction, this field applies to the <b>PostDst</b> – the post destination operand. Encoding: <ul style="list-style-type: none"> <li>00 = “ARF”. Architecture Register File (<b>a#</b>, <b>acc#</b>, <b>f#</b>, <b>n#</b>, <b>null</b>, <b>ip</b>, etc.)</li> <li>01 = “GRF”. General Register File (<b>r#</b>)</li> <li>10 = “MRF”. Message Register File (<b>m#</b>)</li> <li>11 = <b>reserved</b></li> </ul>

The following tables describe the Destination Register Region based on the access mode and addressing mode.



**Table 7-6. Destination Register Region in Direct + Align16 mode**

Bits	Description
15	<b>Dst.AddrMode – Destination Address Mode.</b> This field is the <i>AddrMode</i> for the destination operand. (See section 0 for definition of <i>AddrMode</i> .) For a <i>send</i> instruction, this field applies to <b>PostDst</b> – the post destination operand. Addressing mode for <i>CurrDst</i> (current destination operand) is fixed as Direct. (See Instruction Reference chapter for <i>CurrDst</i> and <i>PostDst</i> .)
14:13	Reserved: MBZ
12:5	<b>Dst.RegNum – Destination Register Number.</b> This field is the <i>RegNum</i> field for the destination operand. (See section 0 for definitions of <i>RegNum</i> .) For a <i>send</i> instruction, this field applies to <b>PostDst</b> .
4	<b>Dst.SubRegNum[4].</b> This is the 16-byte aligned sub-register address. (See section 0 for definitions of <i>SubRegNum</i> ) For a <i>send</i> instruction, this field applies to <b>CurrDst</b> .
3:0	<b>Dst.ChanEn – Destination Channel Enable.</b> The channel enable field for the destination operand. (See section 0 for definitions of <i>ChanEn</i> ) For a <i>send</i> instruction, this field applies to the <b>CurrDst</b> .

**Table 7-7. Destination Register Region in Direct+Align1 mode**

Bits	Description
15	<b>Dst.AddrMode – Destination Address Mode.</b> This field is the <i>AddrMode</i> for the destination operand. For a <i>send</i> instruction, it applies to <b>PostDst</b> . Addressing mode for <b>CurrDst</b> is fixed as Direct.
14:13	<b>Dst.HorzStride – Destination Horizontal Stride.</b> This field is the <i>HorzStride</i> for the destination operand. For a <i>send</i> instruction, this field applies to <b>CurrDst</b> . <b>PostDst</b> only uses the register number.
12:5	<b>Dst.RegNum – Destination Register Number.</b> This field is the <i>RegNum</i> field for the destination operand. For a <i>send</i> instruction, this field applies to <b>PostDst</b> .
4:0	<b>Dst.SubRegNum – Destination Sub-Register Number.</b> This field is the <b>SubRegNum</b> for the destination operand. (See section 0 for definition of <i>SubRegNum</i> ) For a <i>send</i> instruction, this field applies to <b>CurrDst</b> .

**Table 7-8. Destination Register Region in Indirect+Align16 mode**

Bits	Description
15	<b>Dst.AddrMode – Destination Address Mode.</b> This field is the <i>AddrMode</i> for the destination operand. For a <i>send</i> instruction, this field applies to <b>PostDst</b> . Addressing mode for <b>CurrDst</b> is fixed as Direct.
14:13	Reserved: MBZ
12:10	<b>Dst.AddrSubRegNum – Destination Address Sub-Register Number.</b> This field is the <i>AddrSubRegNum</i> for the destination operand. (See section 0 for definition of <i>AddrSubRegNum</i> .) For a <i>send</i> instruction, this field applies to <b>PostDst</b> .



Bits	Description
9:4	<b>Dst.AddrImm[9:4]</b> This is the half-register aligned <i>AddrImm</i> field for the destination operand. (See section 0 for definition of <i>AddrImm</i> ) For a <i>send</i> instruction, this field applies to <b>PostDst</b> .
3:0	<b>Dst.ChanEn – Destination Channel Enable.</b> The channel enable field for the destination operand. For a <i>send</i> instruction, this field applies to the <b>CurrDst</b> .

**Table 7-9. Destination Register Region in Indirect+Align1 mode**

Bits	Description
15	<b>Dst.AddrMode – Destination Address Mode.</b> This field is the <i>AddrMode</i> for the destination operand. For a <i>send</i> instruction, this field applies to <b>PostDst</b> . Addressing mode for <b>CurrDst</b> is fixed as Direct.
14:13	<b>Dst.HorzStride – Destination Horizontal Stride</b> This field is the <i>HorzStride</i> for the destination operand. For a <i>send</i> instruction, this field applies to <b>CurrDst</b> . <b>PostDst</b> only uses the register number.
12:10	<b>Dst.AddrSubRegNum – Destination Address Sub-Register Number.</b> This field is the <i>AddrSubRegNum</i> for the destination operand. For a <i>send</i> instruction, this field applies to <b>PostDst</b> .
9:0	<b>Dst.AddrImm – Destination Address Immediate.</b> This field is the byte-aligned <i>AddrImm</i> for the destination operand. For a <i>send</i> instruction, this field applies to <b>PostDst</b> .

### 7.2.3.2 3-src Instructions

This section describes the field in DW1 of the 3-src instruction format.

**Table 7-10. Instruction DW1**

Bits	Description
31:24	<b>Destination Register Number.</b> This field contains the destination register number.
23:21	<b>Destination Subregister Number.</b> This field contains the destination subregister number.
20:17	<b>Destination Channel Enable.</b> Four channel enables are defined for controlling which channels will be written into the destination region. These channel mask bits are applied in a modulo-four manner to all <i>ExecSize</i> channels. There is 1-bit Channel Enable for each channel within the group of 4. If the bit is cleared, the write for the corresponding channel is disabled. If the bit is set, the write is enabled. Mnemonic for the bit being set for the group of 4 is “x”, “y”, “z”, and “w”, respectively, where “x” corresponds to Channel 0 in the group and “w” corresponds to channel 3 in the group. 0 = Write Disabled 1 = Write Enabled (normal)
16:15	Reserved: MBZ
14:13	Reserved: MBZ
12:10	Reserved: MBZ



Bits	Description
9:8	<b>Source2 Modifier.</b> This field contains the modifier for source2. Refer to Table. 5-5 for the encoding.
7:6	<b>Source1 Modifier.</b> This field contains the modifier for source1. Refer to Table. 5-5 for the encoding.
5:4	<b>Source0 Modifier.</b> This field contains the modifier for source0. Refer to Table. 5-5 for the encoding.
3	Reserved: MBZ
2	Reserved: MBZ
1	<b>Flag Subregister Number.</b> This field contains the flag subregister number for instructions with non-zero Conditional Modifier.
0	<b>Destination Register File.</b> This field contains the register file for destination. Only GRF and MRF are valid. 0 = GRF 1 = MRF

## 7.2.4 Instruction Source-0 Doubleword (DW2)

### 7.2.4.1 1-src and 2-src Instructions

Instruction Source-0 Doubleword (DW2) contains the first source operand and also flag register number.

- Table 7-11 shows the field definition for Direct Addressing with Align16.
- Table 7-12 shows the field definition for Direct Addressing with Align1.
- Table 7-13 shows the field definition for Indirect Addressing with Align16.
-



Table 7-14 shows the field definition for Indirect Addressing with Align1.

**Table 7-11. Instruction Source-0 Doubleword in Direct+Align16 mode**

Bits	Description
31:26	Reserved: MBZ
25	<b>FlagSubRegNum – Flag Sub-Register Number.</b> This field specifies the sub-register number for a flag register operand. There are two sub-registers in the flag register. Each sub-register contains 16 flag bits. The selected flag sub-register is the source for predication if predication is enabled for the instruction. It is the destination to store conditional flag bits if conditional modifier is enabled for the instruction. The same flag sub-register can be both the predication source and conditional destination, if both predication and conditional modifier are enabled.
24:21	<b>Src0.VertStride – Source-0 Vertical Stride.</b> This field is the <i>VertStride</i> for <src0> operand. (See section 0 for definition of <i>VertStride</i> ) It is ignored if <src0> is an immediate operand.
20	Reserved: MBZ
19:16	<b>Src0.ChanSel[7:4]</b> This is bits [7:4] of the <i>ChanSel</i> field for <src0> operand. (See section 0 for definition of <i>ChanSel</i> ). It is ignored if <src0> is an immediate operand.
15	<b>Src0.AddrMode – Source-0 Address Mode.</b> This field is the <i>AddrMode</i> for <src0> operand. (See section 0 for definition of <i>AddrMode</i> ) It is ignored if <src0> is an immediate operand.
14:13	<b>Src0.SrcMod – Source-0 Source Modifier.</b> This field is the <i>SrcMod</i> for source operand <src0>. (See section 0 for definition of <i>SrcMod</i> ) It is ignored if <src0> is an immediate operand.
12:5	<b>Src0.RegNum – Source-0 Register Number</b> This is the <i>RegNum</i> field for source operand <src0>. (See section 0 for definition of <i>RegNum</i> .) It is ignored if <src0> is an immediate operand.
4	<b>Src0.SubRegNum[4]</b> This is the 16-byte aligned sub-register address for source operand <src0>. (See section 0 for definition of <i>SubRegNum</i> ) It is ignored if <src0> is an immediate operand.
3:0	<b>Src0.ChanEn – Source-0 Channel Enable</b> This is the <i>ChanEn</i> field for source operand <src0>. (See section 0 for definitions of <i>ChanEn</i> ) It is ignored if <src0> is an immediate operand.

**Table 7-12. Instruction Source-0 Doubleword in Direct+Align1 mode**

Bits	Description
31:26	Reserved: MBZ
25	<b>FlagSubRegNum – Flag Sub-Register Number.</b> This field specifies the sub-register number for a flag register operand.
24:21	<b>Src0.VertStride – Source-0 Vertical Stride</b> This is the <i>VertStride</i> field for <src0> operand. (See section 0 for definition of <i>VertStride</i> ) It is ignored if <src0> is an immediate operand.



Bits	Description
20:18	<b>Src0.Width</b> . This is the <i>Width</i> field for source operand <src0>. (See section 0 for definition of <i>Width</i> ) It is ignored if <src0> is an immediate operand.
17:16	<b>Src0.HorzStride</b> . This is the <i>HorzStride</i> field for source operand <src0>. (See section 0 for definition of <i>HorzStride</i> ) It is ignored if <src0> is an immediate operand.
15	<b>Src0.AddrMode – Source-0 Address Mode</b> . This is the <i>AddrMode</i> for source operand <src0>. (See section 0 for definition of <i>AddrMode</i> ) It is ignored if <src0> is an immediate operand.
14:13	<b>Src0.SrcMod – Source-0 Source Modifier</b> . This is the <i>SrcMod</i> field for source operand <src0>. (See section 0 for definition of <i>SrcMod</i> ) It is ignored if <src0> is an immediate operand.
12:5	<b>Src0.RegNum – Source-0 Register Number</b> . This is the <i>RegNum</i> field for source operand <src0>. (See section 0 for definition of <i>RegNum</i> .) It is ignored if <src0> is an immediate operand.
4:0	<b>Src0.SubRegNum – Source-0 Sub-Register Number</b> . This is the <i>SubRegNum</i> field for source operand <src0>. (See section 0 for definition of <i>SubRegNum</i> ) It is ignored if <src0> is an immediate operand.

**Table 7-13. Instruction Source-0 Doubleword in Indirect+Align16 mode**

Bits	Description
31:26	Reserved: MBZ
25	<b>FlagSubRegNum – Flag Sub-Register Number</b> . This field specifies the sub-register number for a flag register operand.
24:21	<b>Src0.VertStride – Source-0 Vertical Stride</b> . This is the <i>VertStride</i> field for <src0> operand. (See section 0 for definition of <i>VertStride</i> ) It is ignored if <src0> is an immediate operand.
20	Reserved: MBZ
19:16	<b>Src0.ChanSel[7:4] – Source-0 Channel Select</b> . This is bits [7:4] of the <i>ChanSel</i> field for <src0> operand. (See section 0 for definition of <i>ChanSel</i> .) It is ignored if <src0> is an immediate operand.
15	<b>Src0.AddrMode – Source-0 Address Mode</b> . This is the <i>AddrMode</i> for source operand <src0>. (See section 0 for definition of <i>AddrMode</i> ) It is ignored if <src0> is an immediate operand.
14:13	<b>Src0.SrcMod – Source-0 Source Modifier</b> . This is the <i>SrcMod</i> field for source operand <src0>. (See section 0 for definition of <i>SrcMod</i> ) It is ignored if <src0> is an immediate operand.
12:10	<b>Src0.AddrSubRegNum – Source-0 Address Sub-Register Number</b> . This is the <i>AddrSubRegNum</i> field for source operand <src0>. (See section 0 for definition of <i>AddrSubRegNum</i> .) It is ignored if <src0> is an immediate operand.



Bits	Description
9:4	<b>Src0.AddrImm[9:4] – Source-0 Address Immediate.</b> This contains the half-register aligned <i>AddrImm</i> field ((bits [9:4]) for <src0>. (See section 0 for definition of <i>AddrImm</i> ) It is ignored if <src0> is an immediate operand.
3:0	<b>Src0.ChanEn – Source-0 Channel Enable .</b> This is the <i>ChanEn</i> field for source operand <src0>. (See section 0 for definitions of <i>ChanEn</i> ) It is ignored if <src0> is an immediate operand.



**Table 7-14. Instruction Source-0 Doubleword in Indirect+Align1 mode**

Bits	Description
31:26	Reserved: MBZ
25	<b>FlagSubRegNum – Flag Sub-Register Number.</b> This field specifies the sub-register number for a flag register operand.
24:21	<b>Src0.VertStride – Source-0 Vertical Stride.</b> This is the <i>VertStride</i> field for <src0> operand. (See section 0 for definition of <i>VertStride</i> ) It is ignored if <src0> is an immediate operand.
20:18	<b>Src0.Width.</b> This is the <i>Width</i> field for source operand <src0>. (See section 0 for definition of <i>Width</i> ) It is ignored if <src0> is an immediate operand.
17:16	<b>Src0.HorzStride.</b> This is the <i>HorzStride</i> field for source operand <src0>. (See section 0 for definition of <i>HorzStride</i> ) It is ignored if <src0> is an immediate operand.
15	<b>Src0.AddrMode – Source-0 Address Mode.</b> This is the <i>AddrMode</i> for source operand <src0>. (See section 0 for definition of <i>AddrMode</i> ) It is ignored if <src0> is an immediate operand.
14:13	<b>Src0.SrcMod – Source-0 Source Modifier.</b> This is the <i>SrcMod</i> field for source operand <src0>. (See section 0 for definition of <i>SrcMod</i> ) It is ignored if <src0> is an immediate operand.
12:10	<b>Src0.AddrSubRegNum – Source-0 Address Sub-Register Number.</b> This is the <i>AddrSubRegNum</i> field for source operand <src0>. (See section 0 for definition of <i>AddrSubRegNum</i> .) It is ignored if <src0> is an immediate operand.
9:0	<b>Src0.AddrImm – Source-0 Address Immediate.</b> This is the byte aligned <i>AddrImm</i> field for <src0>. (See section 0 for definition of <i>AddrImm</i> ) It is ignored if <src0> is an immediate operand.

### 7.2.4.2 3-src Instructions

This section describes the field in DW2 and DW3 of the 3-src instruction format.

**Table 7-15. Instruction DW2 and DW3**

DW	Bits	Description
DW3	31:30	Reserved: MBZ
	29:22	<b>Source2 Register Number.</b> This field contains the register number for source2.
	21:19	<b>Source2 Subregister Number.</b> This field contains the subregister number for source2.
	18:11	<b>Source2 Swizzle.</b> This field contains the swizzle control for source2. Refer to Table.5-5 for encoding.
	10:10	<b>Source2 Replication Control.</b> This field controls replication for source2. Refer to Table.5-5 for encoding.
	9:9	Reserved: MBZ
	8:1	<b>Source1 Register Number.</b> This field contains the register number for source1.





DW	Bits	Description
	0	<b>Source1 Subregister Number.</b> This field contains the subregister number for source1.
DW2	31:30	<b>Source1 Subregister Number.</b> This field contains the subregister number for source1.
	29:22	<b>Source1 Swizzle.</b> This field contains the swizzle control for source1. Refer to Table.5-5 for encoding.
	21:21	<b>Source1 Replication Control.</b> This field controls replication for source1. Refer to Table.5-5 for encoding.
	20:20	Reserved: MBZ
	19:12	<b>Source0 Register Number.</b> This field contains the register number for source0.
	11:9	<b>Source0 Subregister Number.</b> This field contains the subregister number for source0.
	8:1	<b>Source0 Swizzle.</b> This field contains the swizzle control for source0. Refer to Table.5-5 for encoding.
	0:0	<b>Source0 Replication Control.</b> This field controls replication for source0. Refer to Table.5-5 for encoding.

## 7.2.5 Instruction Source-1 Doubleword (DW3)

Source-1 Doubleword (DW3) contains the second source operand (<src1>) and is used to hold the 32-bit immediate source (imm32 as <src0> or <src1>). Table 7-16 and Table 7-17 define the fields in this doubleword with the following exceptions:

- If <src0> is an immediate operand, this doubleword contains **imm32** for <src0>.
- If <src1> is an immediate operand, this doubleword contains **imm32** for <src1>.
- If the instruction is a send, bit 31 of this doubleword contains **EOT** field.
  - If <src1> is immediate, the remaining 31 bits in this doubleword is **MsgDescpt31**.
  - If <src1> is a register, <src1> must be a0.0. The rest of this doubleword will be configured accordingly.
- If indirect address is supported for <src1>, Table 7-18 and Table 7-19 define the fields in DW3 for indirectly addressed <src1> in Align16 and Align1 modes.

**Table 7-16. Instruction Source-1 Doubleword in Direct + Align16 mode**

Bits	Description
31:25	Reserved: MBZ
24:21	<b>Src1.VertStride – Source-1 Vertical Stride.</b> This field is the <i>VertStride</i> for <src1> operand. (See section 0 for definition of <i>VertStride</i> ) It is ignored if <src1> is an immediate operand.
20	Reserved: MBZ



Bits	Description
19:16	<b>Src1.ChanSel[7:4]</b> This contains bits [7:6] of the <i>ChanSel</i> field for <src1> operand. (See section 0 for definition of <i>ChanSel</i> ) It is ignored if <src1> is an immediate operand.
15	Reserved: MBZ
14:13	<b>Src1.SrcMod – Source-1 Source Modifier.</b> This field is the <i>SrcMod</i> for <src1> operand. (See section 0 for definition of <i>SrcMod</i> ) It is ignored if <src1> is an immediate operand.
12:5	<b>Src1.RegNum.</b> This field is the <i>RegNum</i> field for <src1> operand. (See section 0 for definition of <i>RegNum</i> .) It is ignored if <src1> is an immediate operand.
4	<b>Src1.SubRegNum[4].</b> This field is bit [4] of the <i>SubRegNum</i> field for <src1>. (See section 0 for definition of <i>SubRegNum</i> ) It is ignored if <src1> is an immediate operand.
3:0	<b>Src1.ChanEn – Source-1 Channel Enable.</b> It is the channel enable field for <src1>. (See section 0 for definitions of <i>ChanEn</i> ) It is ignored if <src1> is an immediate operand.

**Table 7-17. Instruction Source-1 Doubleword in Direct + Align1 mode**

Bits	Description
31:25	Reserved: MBZ
24:21	<b>Src1.VertStride – Source-1 Vertical Stride.</b> This field is the <i>VertStride</i> for <src1> operand. (See section 0 for definition of <i>VertStride</i> ) It is ignored if <src1> is an immediate operand.
20:18	<b>Src1.Width.</b> This is the <i>Width</i> field for source operand <src1>. (See section 0 for definition of <i>Width</i> ) It is ignored if <src1> is an immediate operand.
17:16	<b>Src1.HorzStride.</b> This is the <i>HorzStride</i> field for source operand <src1>. (See section 0 for definition of <i>HorzStride</i> ) It is ignored if <src1> is an immediate operand.
15	Reserved: MBZ
14:13	<b>Src1.SrcMod – Source-1 Source Modifier.</b> This field is the <i>SrcMod</i> for <src1> operand. (See section 0 for definition of <i>SrcMod</i> ) It is ignored if <src1> is an immediate operand.
12:5	<b>Src1.RegNum – Source-1 Register Number.</b> This is the <i>RegNum</i> field for source operand <src1>. (See section 0 for definition of <i>RegNum</i> .) It is ignored if <src1> is an immediate operand.
4:0	<b>Src1.SubRegNum – Source-1 Sub-Register Number.</b> This is the <i>SubRegNum</i> field for source operand <src1>. (See section 0 for definition of <i>SubRegNum</i> ) It is ignored if <src1> is an immediate operand.



**Table 7-18. Instruction Source-1 Doubleword in Indirect+Align16 mode**

Bits	Description
31:25	Reserved: MBZ
24:21	<b>Src1.VertStride – Source-1 Vertical Stride</b> This is the <i>VertStride</i> field for <src1> operand. (See section 0 for definition of <i>VertStride</i> ) It is ignored if <src1> is an immediate operand.
20	Reserved: MBZ
19:16	<b>Src1.ChanSel[7:4] – Source-1 Channel Select</b> This is bits [7:4] of the <i>ChanSel</i> field for <src1> operand. (See section 0 for definition of <i>ChanSel</i> ). It is ignored if <src1> is an immediate operand.
15	<b>Src1.AddrMode – Source-1 Address Mode</b> This is the <i>AddrMode</i> for source operand <src1>. (See section 0 for definition of <i>AddrMode</i> ) It is ignored if <src1> is an immediate operand.
14:13	<b>Src1.SrcMod – Source-1 Source Modifier</b> This is the <i>SrcMod</i> field for source operand <src1>. (See section 0 for definition of <i>SrcMod</i> ) It is ignored if <src1> is an immediate operand.
12:10	<b>Src1.AddrSubRegNum – Source-1 Address Sub-Register Number</b> This is the <i>AddrSubRegNum</i> field for source operand <src1>. (See section 0 for definition of <i>AddrSubRegNum</i> .) It is ignored if <src1> is an immediate operand.
9:4	<b>Src1.AddrImm[9:4] – Source-1 Address Immediate</b> This contains the half-register aligned <i>AddrImm</i> field ((bits [9:4]) for <src1>. (See section 0 for definition of <i>AddrImm</i> ) It is ignored if <src1> is an immediate operand.
3:0	<b>Src1.ChanEn – Source-1 Channel Enable</b> This is the <i>ChanEn</i> field for source operand <src1>. (See section 0 for definitions of <i>ChanEn</i> ) It is ignored if <src1> is an immediate operand.

**Table 7-19. Instruction Source-1 Doubleword in Indirect+Align1 mode**

Bits	Description
31:25	Reserved: MBZ
24:21	<b>Src1.VertStride – Source-1 Vertical Stride</b> This is the <i>VertStride</i> field for <src1> operand. (See section 0 for definition of <i>VertStride</i> ) It is ignored if <src1> is an immediate operand.
20:18	<b>Src1.Width</b> This is the <i>Width</i> field for source operand <src1>. (See section 0 for definition of <i>Width</i> ) It is ignored if <src1> is an immediate operand.
17:16	<b>Src1.HorzStride</b> This is the <i>HorzStride</i> field for source operand <src1>. (See section 0 for definition of <i>HorzStride</i> ) It is ignored if <src1> is an immediate operand.



Bits	Description
15	<b>Src1.AddrMode – Source-1 Address Mode</b> This is the <i>AddrMode</i> for source operand <src1>. (See section 0 for definition of <i>AddrMode</i> ) It is ignored if <src1> is an immediate operand.
14:13	<b>Src1.SrcMod – Source-1 Source Modifier</b> This is the <i>SrcMod</i> field for source operand <src1>. (See section 0 for definition of <i>SrcMod</i> ) It is ignored if <src1> is an immediate operand.
12:10	<b>Src1.AddrSubRegNum – Source-1 Address Sub-Register Number</b> This is the <i>AddrSubRegNum</i> field for source operand <src1>. (See section 0 for definition of <i>AddrSubRegNum</i> .) It is ignored if <src1> is an immediate operand.
9:0	<b>Src1.AddrImm – Source-1 Address Immediate</b> This is the byte aligned <i>AddrImm</i> field for <src1>. (See section 0 for definition of <i>AddrImm</i> ) It is ignored if <src1> is an immediate operand.

## 7.3 Opcode Encoding

Byte 0 of the 128-bit instruction word contains the opcode. The opcode uses 7 bits. Bit location 7 in byte 0 is reserved for future opcode extension.

There are total of 48 opcodes defined. These opcodes are encoded and organized into five groups based on the type of operations: Special instructions, move/logic instructions (opcode=00xxxxb), flow control instructions (opcode=010xxxxb), miscellaneous instructions (opcode=011xxxxb), parallel arithmetic instructions (opcode=100xxxxb), and vector arithmetic instructions (opcode=101xxxxb). Opcodes 110xxxxb are reserved.

### 7.3.1 Move and Logic Instructions

This instruction group has an opcode format of 00xxxxxb.

- The opcodes for move instructions (***mov***, ***sel*** and ***movi***) share the common 5 MSBs in the form of 00000xxb.
- The opcodes for logic instructions (***not***, ***and***, ***or***, and ***xor***) share the common 5 MSBs in the form of 00001xxb.
- The opcodes for shift instructions (***shr***, ***shl***, and ***asr***) share the common 4 MSBs in the form of 0001xxb. Bit 2 indicates arithmetic or logic shift (0 = logic, 1 = arithmetic). Bit 1 is always 0 (which is reserved for future extension to support rotation shift as 0 = shift, 1 = rotate). Bit 0 indicates the shift direction (0 = right, 1 = left).
- The opcodes for compare instructions (***cmp*** and ***cmpn***) share the common 6 MSBs in the form of 001000xb. Bit 0 indicates whether it is a normal compare, *cmp*, or a special compare-NaN, *cmpn*.
- This group of instructions does not implicitly update the accumulators.
- Instruction compression applies to this group.



**Table 7-20. Move and Logic Instructions**

Opcode		Instruction	Description	#src	#dst
dec	hex				
1	0x01	<b>mov</b>	Component-wise move	1	1
2	0x02	<b>sel</b>	Component-wise selective move based on predication	2	1
3	0x03	<b>movi</b>	Fast component-wise indexed move	1	1
4	0x04	<b>not</b>	Component-wise one's compliment (bitwise not)	1	1
5	0x05	<b>and</b>	Component-wise logical AND (bitwise and)	2	1
6	0x06	<b>or</b>	Component-wise logical OR (bitwise or)	2	1
7	0x07	<b>xor</b>	Component-wise logical XOR (bitwise xor)	2	1
8	0x08	<b>shr</b>	Component-wise logical shift right	2	1
9	0x09	<b>shl</b>	Component-wise logical shift left	2	1
<b>10</b>	<b>0x0A</b>	<b>dim</b>	<b>Double Precision Floaitng Point Immediate Data Move</b>	<b>1</b>	<b>1</b>
11	0x0A-0x0B	<i>Reserved</i>			
12	0x0C	<b>asr</b>	Component-wise arithmetic shift right	2	1
13-15	0x0D-0x0F	<i>Reserved</i>			
16	0x10	<b>cmp</b>	Component-wise compare, store condition code in destination	2	1
17	0x11	<b>cmpn</b>	Component-wise compare-NaN, store condition code in destination	2	1
18	0x12	<i>Reserved</i>			
19	0x13	<b>Reserved</b>			
20	0x14	<b>Reserved</b>			
21-22	0x15-0x16	<i>Reserved</i>			
23	0x17	<b>Reserved</b>			
24	0x18	<b>Reserved</b>			
25	0x19	<b>Reserved</b>			
26	0x1A	<b>Reserved</b>			
20-31	0x12-0x1F	<i>Reserved</i>			



### 7.3.2 Flow Control Instructions

This instruction group has an opcode format of 010xxxxb.

- This group of instructions does not implicitly update the accumulators.
- Instruction compression is not allowed for this group.

**Table 7-21. Flow Control Instructions**

Opcode		Instruction	Description	#src	#dst
dec	hex				
32	0x20	<b><i>jmp</i></b>	Jump indexed	1	0
33	0x21	<b><i>Reserved</i></b>			
34	0x22	<b><i>if</i></b>	If	0/2	0
35	0x23	<b><i>Reserved</i></b>			
36	0x24	<b><i>else</i></b>	Else	1	0
37	0x25	<b><i>endif</i></b>	End if	0	0
38	0x26	<b><i>case</i></b>	Case – Inside Switch block	0/2	0
39	0x27	<b><i>while</i></b>	While	1	0
40	0x28	<b><i>break</i></b>	Break	1	0
41	0x29	<b><i>cont</i></b>	Continue	1	0
42	0x2A	<b><i>halt</i></b>	Halt	1	0
43	0x2B	<b><i>Reserved</i></b>			
44	0x2C	<b><i>call</i></b>	Subroutine call	1	1
45	0x2D	<b><i>return</i></b>	Subroutine return	1	1
46	0x2E	<b><i>fork</i></b>	go into 16pixel execution mode from 32pixel execution mode	1	0
47	0x2F	<b><i>Reserved</i></b>			

### 7.3.3 Miscellaneous Instructions

This instruction group has an opcode format of 011xxxxb.

- This group of instructions does not implicitly update the accumulators.
- Instruction compression is not allowed for this group.



**Table 7-22. Miscellaneous Instructions**

Opcode		Instruction	Description	#src	#dst
dec	hex				
48	0x30	<b>wait</b>	Wait for (external) notification	1	0
49	0x31	<b>send</b>	Send	1	1
50	0x32	<b>sendc</b>	Conditional Send (based on TDR)	1	1
51-55	0x33-0x37	<i>Reserved</i>			
56	0x38	<b>math</b>	Math functions for extended math pipeline	1/2	1/2
57-63	0x39-0x3F	<i>Reserved</i>			

### 7.3.4 Parallel Arithmetic Instructions

This instruction group has an opcode format of 100xxxxb.

- The opcode for round instructions (**rndu**, **rndd**, **rnde**, and **rndz**) share the common 5 MSBs in the form of 10001xxb, with the lower 2 bits indicate the type of round.
- These instructions implicitly update the accumulators if the Accumulator Disable bit in control register cr0.0 not set.
  - Some instructions such as **frc**, **lzd**, etc, perform the operation after the accumulator. Therefore, when the accumulator is implicitly updated, the content is undefined. Details can be found in ISA Reference Chapter.
- Instruction compression applies to this group.

**Table 7-23. Parallel Arithmetic Instructions**

Opcode		Instruction	Description	#src	#dst
dec	hex				
64	0x40	<b>add</b>	Component-wise addition	2	1
65	0x41	<b>mul</b>	Component-wise multiply	2	1
66	0x42	<b>avg</b>	Component-wise average of the two source operands	2	1
67	0x43	<b>frc</b>	Component-wise floating point truncate-to-minus-infinity fraction	1	1
68	0x44	<b>rndu</b>	Component-wise floating point rounding up (ceiling)	1	1
69	0x45	<b>rndd</b>	Component-wise floating point rounding down (floor)	1	1
70	0x46	<b>rnde</b>	Component-wise floating point rounding toward nearest even	1	1



Opcode		Instruction	Description	#src	#dst
dec	hex				
71	0x47	<b><i>rndz</i></b>	Component-wise floating point rounding toward zero	1	1
72	0x48	<b><i>mac</i></b>	Component-wise multiply accumulate	2	1
73	0x49	<b><i>mach</i></b>	multiply accumulate high	2	1
74	0x4A	<b><i>lzd</i></b>	leading zero detection	1	1
75	0x4B	<b><i>Reserved</i></b>			
76	0x4C	<b><i>Reserved</i></b>			
77	0x4D	<b><i>Reserved</i></b>			
78	0x4E	<b><i>Reserved</i></b>			
79	0x4F	<b><i>Reserved</i></b>			

### 7.3.5 Vector Arithmetic Instructions

- This instruction group has an opcode format of 101xxxxb.
- These instructions implicitly update the accumulators if the Accumulator Disable bit in control register cr0.0 not set.
  - Some instructions such as ***dp4-dp2***, etc, perform the operation after the accumulator. Therefore, when the accumulator is implicitly updated, the content is undefined. Details can be found in ISA Reference Chapter.
- Instruction compression applies to this group.

**Table 7-24. Vector Arithmetic Instructions**

Opcode		Instruction	Description	#src	#dst
dec	hex				
80	0x50	<b><i>sad2</i></b>	2-wide sum of absolute difference	2	1
81	0x51	<b><i>sada2</i></b>	2-wide sad accumulate	2	1
82-83	0x52-0x53	<i>reserved</i>			
84	0x54	<b><i>dp4</i></b>	4-wide dot product for 4-vector	2	1
85	0x55	<b><i>dph</i></b>	4-wide homogenous dot product for 4-vector	2	1
86	0x56	<b><i>dp3</i></b>	3-wide dot product for 4-vector	2	1
87	0x57	<b><i>dp2</i></b>	2-wide dot product for 4-vector	2	1
88	0x58	<i>reserved</i>			





Opcode		Instruction	Description	#src	#dst
dec	hex				
89	0x59	<b>line</b>	Component-wise line equation computation (a multiply-add)	2	1
90	0x5A	<b>pln</b>	Component-wise floating point plane equation computation (a multiply-multiply-add)	2(3)	1
91	0x5B	<b>mad</b>	Component-wise floating point mad computation (a multiple-add)	3	1
92	0x5C	<b>lrp</b>	Component-wise floating point lrp computation (blend)	3	1
93-95	0x5D-0x5F	<i>reserved</i>			

### 7.3.6 Special Instructions

There are two special instructions, namely, *nop* (opcode = 0x7E) and *illegal* (opcode = 0x00).

- *Nop* instruction may be used for instruction padding in memory between two normal instructions to force alignment or to introduce instruction execution delay. Currently, there is no need for between-instruction padding.
- *Illegal* instruction may be used for instruction padding in memory outside the normal instruction sequence such as before or after the kernel program as well as between subroutines.
- *Nop* and *illegal* instructions do not have source operands or destination operand. Therefore, they do not implicitly update the accumulator register. They cannot be compressed.

**Table 7-25. Special Instructions**

Opcode		Instruction	Description	#src	#dst
dec	hex				
0	0x00	<b>illegal</b>	Illegal instruction	0	0
96-124	0x60-0x7C	<i>Reserved</i>			
125	0x7D	<b>nenop</b>	Non-executed No-op	0	0
126	0x7E	<b>nop</b>	No-op	0	0
127	0x7F	<i>Reserved</i>	(may be used as an extension code)		



## 7.4 Native Instruction BNF

The Backus-Naur Form (BNF) grammar identifies the assembly language syntax, which is native to the hardware. It does not include intelligent defaults, assembler pragmas, etc.

### 7.4.1 Instruction Groups

<Instruction>	::=	<UnaryInstruction>   <BinaryAcclInstruction>   <BinaryInstruction>   <TriInstruction>   <JumpInstruction>   <BranchLoopInstruction>   <ElseInstruction>   <BreakInstruction>   <MaskControlInstruction>   <SyncInstruction>   <SpecialInstruction>
<UnaryInstruction>	::=	<Predicate> <UnaryInst> <ExecSize> <Dst> <SrcAcclmm> <InstOptions>
<UnaryInst>	::=	<UnaryOp> <ConditionalModifier> <Saturate>
<UnaryOp>	::=	<b>"mov"   "frc"   "rndu"   "rnda"   "rnde"   "rndz"   "not"   "lzd"</b>
<BinaryInstruction>	::=	<Predicate> <BinaryInst> <ExecSize> <Dst> <Src> <SrcImm> <InstOptions>
<BinaryInst>	::=	<BinaryOp> <ConditionalModifier> <Saturate>
<BinaryOp>	::=	<b>"mul"   "mac"   "mach"   "line"</b>  <b>  "sad2"   "sada2"   "dp4"   "dph"   "dp3"   "dp2"</b>
<BinaryAcclInstruction>	::=	<Predicate> <BinaryAcclInst> <ExecSize> <Dst> <SrcAcc> <SrcImm> <InstrOptions>
<BinaryAcclInst>	::=	<BinaryAccOp> <ConditionalModifier> <Saturate>
<BinaryAccOp>	::=	<b>"avg"   "add"   "sel"</b>  <b>  "and"   "or"   "xor"</b>  <b>  "shr"   "shl"   "asr"</b>  <b>  "cmp"   "cmpn"</b>
<TriInstruction>	::=	<Predicate> <TriInst> <ExecSize> <PostDst> <CurrDst> <TriSrc> <MsgDesc> <InstOptions>



<TriInst> ::= <TriOp> <ConditionalModifier> <Saturate>  
<TriOp> ::= **“send”**

<JumpInstruction> ::= <JumpOp> <RelativeLocation2>  
<JumpOp> ::= **“jmp”**

<BranchLoopInstruction> ::= <Predicate> <BranchLoopOp> <RelativeLocation>  
<BranchLoopOp> ::= **“if” | “iff” | “while”**

<ElseInstruction> ::= <ElseOp> <RelativeLocation>  
<ElseOp> ::= **“else”**

<BreakInstruction> ::= <Predicate> <BreakOp> <LocationStackCtrl>  
<BreakOp> ::= **“break” | “cont” | “halt”**

<SyncInstruction> ::= <Predicate> <SyncOp> <NotifyReg>  
<SyncOp> ::= **“wait”**

<SpecialInstruction> ::= **“do” | “endif” | “nop” | “illegal”**

## 7.4.2 Destination Register

<Dst> ::= <DstOperand>  
| <DstOperandEx>

<DstOperand> ::= <DstReg> <DstRegion> <WriteMask> <DstType>

<DstOperandEx> ::= <AccReg> <DstRegion> <DstType>  
| <FlagReg> <DstRegion> <DstType>  
| <AddrReg> <DstRegion> <DstType>  
| <MaskReg> <DstRegion> <DstType>  
| <MaskStackReg>



| <ControlReg>  
| <IPReg>  
| <NullReg>

<DstReg> ::= <DirectGenReg> | <IndirectGenReg>  
| <DirectMsgReg> | <IndirectMsgReg>

<PostDst> ::= <PostDstReg> <DstRegion> <WriteMask> <DstType>  
| <NullReg>

<PostDstReg> ::= <DirectGenReg> | <IndirectGenReg>

<CurrDst> ::= <DirectAlignedMsgReg>

### 7.4.3 Source Register

#### Source with Accumulator Access and with Immediate

<SrcAcclmm> ::= <SrcAcc>  
| <Imm32> <SrcImmType>

<SrcAcc> ::= <DirectSrcAccOperand>  
| <IndirectSrcOperand>

<DirectSrcAccOperand> ::= <DirectSrcOperand>  
| <SrcArcOperandEx>  
| <AccReg> <SrcType>

<SrcArcOperandEx> ::= <FlagReg> <Region> <SrcType>  
| <AddrReg> <Region> <SrcType>  
| <MaskReg> <Region> <SrcType> [Pre-DevSNB]  
| <MaskStackReg> [Pre-DevSNB]  
| <ControlReg>



| <StateReg>  
| <NotifyReg>  
| <IPReg>  
| <NullReg>

<IndirectSrcOperand> ::= <SrcModifier> <IndirectGenReg> <IndirectRegion> <Swizzle > <SrcType>

### Source without Accumulator Access

<Src> ::= <DirectSrcOperand>  
| <IndirectSrcOperand>

< DirectSrcOperand > ::= <SrcModifier> <DirectGenReg> <Region> <Swizzle> <SrcType>  
| <SrcArcOperandEx>

<TriSrc> ::= <SrcModifier> <DirectGenReg> <Region> <Swizzle> <SrcType>  
| <NullReg>

<MsgDesc> ::= <ImmDesc>  
| <Reg32>

<Reg32> ::= <DirectGenReg> <Region> <SrcType>

### Source without Accumulator Access or IP Access

<SrcImm> ::= <DirectSrcOperand>  
| <Imm32> <SrcImmType>

## 7.4.4 Address Registers

<AddrParam> ::= <AddrReg> <ImmAddrOffset>

<ImmAddrOffset> ::= “”  
| “,” <ImmAddrNum>



## 7.4.5 Register Files and Register Numbers

<DirectGenReg>	::= <GenRegFile> <GenRegNum> <GenSubRegNum>
<IndirectGenReg>	::= <GenRegFile> “[ <AddrParam> ]”
<GenRegFile>	::= “r”
<GenRegNum>	::= “0”...“127”
<GenSubRegNum>	::= “”   “.0”...“.7”   “.0”...“.15”   “.0”...“.31”
<DirectMsgReg>	::= <DirectAlignedMsgReg> <MsgSubRegNum>
<DirectAlignedMsgReg>	::= <MsgRegFile> <MsgRegNum>
<IndirectMsgReg>	::= <MsgRegFile> “[ <AddrParam> ]”
<MsgRegFile>	::= “m”
<MsgRegNum>	::= “0”...“15”
<MsgSubRegNum>	::= <GenSubRegNum>
<AddrReg>	::= <AddrRegFile> <AddrSubRegNum>
<AddrRegFile>	::= “a0”
<AddrSubRegNum>	::= “”   “.0”...“.7”
<AccReg>	::= “acc” <AccRegNum><AccSubRegNum>
<AccRegNum>	::= “0”   “1”
<AccSubRegNum>	::= <GenSubRegNum>
<FlagReg>	::= “f0” <FlagSubRegNum>
<FlagSubRegNum>	::= “”   “.0”...“.1”



<NotifyReg> ::= "n" <NotifyRegNum>  
<NotifyRegNum> ::= "0" ... "2"  
  
<StateReg> ::= "sr0" <StateSubRegNum>  
<StateSubRegNum> ::= ".0" ... ".1"  
  
<ControlReg> ::= "cr0" <ControlSubRegNum>  
<ControlSubRegNum> ::= ".0" ... ".2"  
  
<IPReg> ::= "ip"  
  
<NullReg> ::= "null"

## 7.4.6 Relative Location and Stack Control

<RelativeLocation> ::= <imm16>  
<RelativeLocation2> ::= <imm32> | <reg32>  
<LocationStackCtrl> ::= <imm32>

## 7.4.7 Regions

<DstRegion> ::= "<" <HorzStride> ">"  
  
<IndirectRegion> ::= <Region> | <RegionWH> | <RegionV>  
  
<Region> ::= "<" <VertStride> " ," <Width> " ," <HorzStride> ">"  
<RegionWH> ::= "<" <Width> " ," <HorzStride> ">"  
<RegionV> ::= "<" <VertStride> ">"  
  
<VertStride> ::= "0" | "1" | "2" | "4" | "8" | "16" | "32"  
<Width> ::= "1" | "2" | "4" | "8" | "16"



<HorzStride> ::= "0" | "1" | "2" | "4"

## 7.4.8 Types

<SrcType> ::= ".f" | ".ud" | ".d" | ".uw" | ".w" | ".ub" | ".b"

<SrcImmType> ::= <SrcType> | ".v" | ".vf"

<DstType> ::= <SrcType>

## 7.4.9 Write Mask

<WriteMask> ::= ""  
| ".x" | ".y" | ".z" | ".w"  
| ".xy" | ".xz" | ".xw" | ".yz" | ".yw" | ".zw"  
| ".xyz" | ".xyw" | ".xzw" | ".yzw"  
| ".xyzw"

## 7.4.10 Swizzle Control

<Swizzle> ::= ""  
| ".<ChanSel>  
| ".<ChanSel> <ChanSel> <ChanSel> <ChanSel>

<ChanSel> ::= "x" | "y" | "z" | "w"

## 7.4.11 Immediate Values

<ImmAddrNum> ::= "-512" ... "511"

<Imm32> ::= "0" ... " $\pm 1.0 \cdot 2^{-128} \dots 127$ " | "0" ... " $2^{32}-1$ " | " $-2^{31}$ " ... " $2^{31}-1$ "

<Imm16> ::= "0" ... " $2^{16}-1$ " | " $-2^{15}$ " ... " $2^{15}-1$ "

<ImmDesc> ::= "0" ... " $2^{32}-1$ "





## 7.4.12 Predication and Modifiers

### Instruction Predication

<Predicate> ::= “  
| “(” <PredState> <FlagReg> <PredCntrl> “)”

<PredState> ::= “  
| “+”  
| “-”

<PredCntrl> ::= “  
| “.x” | “.y” | “.z” | “.w”  
| “.any2h” | “.all2h”  
| “.any4h” | “.all4h”  
| “.any8h” | “.all8h”  
| “.any16h” | “.all16h”  
| “.anyv” | “.allv”

### Source Modification

<SrcModifier> ::= “  
| “-”  
| “(abs)”  
| “-” “(abs)”

### Instruction Modification

<ConditionalModifier> ::= “  
| <CondMod> “.” <FlagReg>

<CondMod> ::= “.z” | “.e” | “.nz” | “.ne” | “.g” | “.ge” | “.l” | “.le” | “.o” | “.r” | “.u”

<Saturate> ::= “  
| “.sat”



## Execution Size

<ExecSize> ::= “(“ <NumChannels> “)”  
<NumChannels> ::= “1” | “2” | “4” | “8” | “16” | “32”

## 7.4.13 Instruction Options

<InstOptions> ::= “”  
| “{” <InstOption> “}”  
| “{” <InstOption> <InstOptionEx> “}”

<InstOptionEx> ::= “”  
| “,” <InstOption> <InstOptionEx>

<InstOption> ::= <AccessMode>  
| <ComprCtrl>  
| <ThreadCtrl>  
| <DependencyCtrl>  
| <MaskCtrl> [Pre-DevSNB]  
| <SendCtrl>  
| <AccWrCtrl> [DevSNB+]  
| <WECtrl> [DevSNB+]

<AccessMode> ::= “Align1” | “Align16”

<ComprCtrl> ::= “SecHalf” | “Compr”

<ThreadCtrl> ::= “Switch”

<DependencyCtrl> ::= “NoDDChk” | “NoDDClr”

<MaskCtrl> ::= “NoMask”

<SendCtrl> ::= “EOT”

*Note for Assembler: Compression control “Compr” has a direct map to the binary instruction word. It may be omitted as long as the Assembler is able to determine whether an instruction is in compressed mode or not based on the execution size and the mode of operation.*



## 7.5 Deprecated Features

### 7.5.1 Defeatured Instructions

The following instructions are removed from implementation mainly due to implementation cost/schedule reasons. They are candidates for future generations.

- Sum of Absolute Difference 4 (sad4)
- Sum of Absolute Difference Accumulate 4 (sada4)
- Add Accumulate (aac)
- Min (min)
- Max (max)
- Next (next)
- Swizzle (swz)
- Dot Product Accumulate 2 (dpa2)
- Rotation Shift Left (rsl)
- Rotation Shift Right (rsr)

### 7.5.2 Others

The following features are also deprecated from implementation.

- Restricted 16-bit Half Floating-Point Numbers



## 8. Instruction Set Reference

This chapter describes the functions of instructions. Each instruction is given a different page and the pages are sorted in alphabetical order according to assembly language mnemonic.

### 8.1 Conventions

#### 8.1.1 Pseudo Code Format

The instructions are explained in the following pseudo-code format that resembles the assembly instruction format.

```
[(<pred>)] opcode (<exec_size>) <dst> <src0> [<src1>]
```

Square bracket “[ ]” is used to signify that the field is optional. Saturation modifier and instruction options are omitted for simplicity.

#### 8.1.2 General Macros and Definitions

*INST\_BYTE\_COUNT* is defined as a constant of 16 bytes.

```
#define INST_BYTE_COUNT    16    // byte count of instruction word
```

Function *floor()* converts a floating point value to an integral floating point value. For a given floating point value, from its closest two integral float values, function *floor()* returns the one that is closer to the negative infinity. For example, *floor(1.3f) = 1.0f*, and *floor(-1.3f) = -2.0f*.

```
float floor(float g)
{
    return maximum( any integral float f: f <= g)
}
```

Function *Condition()* takes the conditional signals {SN, ZR, OF, IN, NC} of *result*, generates a Boolean data according to a conditional evaluation controlled by the conditional modifier *cmod*, and returns the Boolean data.

```
Bool Condition(result, cmod)
{
}
```

Function *ConditionNaN()* takes the conditional signals {SN, ZR, OF, IN, NC, NS} of *result*, generates a Boolean data according to a conditional evaluation controlled by the conditional modifier *cmod*, and



returns the Boolean data. The only difference between *Condition()* and *ConditionNaN()* is that *ConditionNaN()* uses the NS (NaN of the second source) signal.

```
Bool ConditionNaN(result, cmod)
{
}
```

Function *Jump()* jumps the instruction sequence from the current instruction location by *InstCount* number of instructions. If *InstCount* is a positive number, it jumps forward; if *InstCount* is a negative number, it jumps backward; if *InstCount* is zero, it is effectively an infinite loop on the current instruction.

```
void Jump(int InstCount)
{
    IP = IP + (InstCount * INST_BYTE_COUNT)
}
```

## 8.2 Evaluate Write Enable

The WrEn should be evaluated as below.

```
if (WECtrl == 1) {
    for (n = 0; n < exec_size; n++) {
        WrEn[n] = 1;
    }
} else {
    for (n = 0; n < exec_size; n++) {
        if (PclIP[n] == ExIP) {
            WrEn[n] = 1;
        } else {
            WrEn[n] = 0;
        }
    }
}

if (PredCtrl != "0000") {
    for (n = 0; n < exec_size; n++) {
        WrEn[n] = WrEn[n] & PMask[n];
    }
}
```



```

for (n = exec_size; n < 32; n++) {
    WrEn[n] = 0;
}

```

## 8.3 Instruction Description

The rest of the chapter contains the description of instructions.

### 8.3.1 add – Addition

Opcode	Instruction	Description
64 (0x40)	add <dst> <src0> <src1>	Component-wise addition of <src0> and <src1> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]
•	•	•	•	[INT]	[INT]
•	•	•	•	[INT]	[FLT]

#### Format:

```
[[<pred>]] add[.<cmod>] (<exec_size>) <dst> <src0> <src1>
```

#### Syntax:

```
[[<pred>]] add[.<cmod>] (<exec_size>) reg reg reg
[[<pred>]] add[.<cmod>] (<exec_size>) reg reg imm32
```

#### Pseudocode:

```

Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = src0.chan[n] + src1.chan[n];
    }
}

```

#### Description:

The *add* instruction performs component-wise addition of <src0> and <src1> and stores the results in <dst>.



Addition of two floating point numbers follows rules in Table 8-1 (or



Table 8-2), if the current floating point mode is IEEE mode (or ALT mode).

**Table 8-1. Floating point addition of A (column) and B (row) in IEEE mode**

	<b>-inf</b>	<b>-finite</b>	<b>-denorm</b>	<b>-0</b>	<b>+0</b>	<b>+denorm</b>	<b>+finite</b>	<b>+inf</b>	<b>NaN</b>
<b>-inf</b>	-inf	-inf	-inf	-inf	-inf	-inf	-inf	NaN	NaN
<b>-finite</b>	-inf	*	A	A	A	A	**	+inf	NaN
<b>-denorm</b>	-inf	B	-0	-0	+0	+0	B	+inf	NaN
<b>-0</b>	-inf	B	-0	-0	+0	+0	B	+inf	NaN
<b>+0</b>	-inf	B	+0	+0	+0	+0	B	+inf	NaN
<b>+denorm</b>	-inf	B	+0	+0	+0	+0	B	+inf	NaN
<b>+finite</b>	-inf	**	A	A	A	A	***	+inf	NaN
<b>+inf</b>	NaN	+inf	+inf	+inf	+inf	+inf	+inf	+inf	NaN
<b>NaN</b>	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN

Notes:

- \* Result can be { -finite}
- \*\* Result can be { -finite, -0, +0, +finite}
- \*\*\* Result can be { +finite}





**Table 8-2. Floating point addition of A (column) and B (row) in ALT mode**

	- fmax	-finite	-denorm	-0	+0	+denorm	+finite	+ fmax	****
-fmax	-fmax	-fmax	-fmax	-fmax	-fmax	-fmax	-finite	+0	
-finite	-fmax	*	A	A	A	A	**	+fmax	
-denorm	-fmax	B	-0	-0	+0	+0	B	+fmax	
-0	-fmax	B	-0	-0	+0	+0	B	+fmax	
+0	-fmax	B	+0	+0	+0	+0	B	+fmax	
+denorm	-fmax	B	+0	+0	+0	+0	B	+fmax	
+finite	-finite	**	A	A	A	A	***	+fmax	
+fmax	+0	+fmax	+fmax	+fmax	+fmax	+fmax	+fmax	+fmax	
****									

Notes:

- \* Result can be { -fmax, -finite}
- \*\* Result can be { -finite, -0, +0, +finite}
- \*\*\* Result can be { +fmax, +finite}
- \*\*\*\* Result is undefined If any of A and/or is {-inf, +inf, NaN}

### 8.3.2 and – Logical And

Opcode	Instruction	Description
5 (0x05)	and <dst> <src0> <src1>	Performing component-wise logic AND of <src0> and <src1> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•		•	•	[INT]	[INT]

**Format:**

[(<pred>)] and[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] and[.<cmod>] (<exec\_size>) reg reg reg  
 [(<pred>)] and[.<cmod>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = src0.chan[n] & src1.chan[n];
    }
}
```

**Description:**

The *and* instruction performs component-wise logic AND operation between <src0> and <src1> and stores the results in <dst>. Source modifiers are allowed.

Accumulator register is allowed to be the destination of this instruction with the restrictions listed below.

**Restrictions:**

Sign (SN) and Overflow (OF) conditions are undefined for this logic instruction. Consequently, saturation modifier (.sat) is not allowed.

This instruction only applies to integer data types. The behavior is undefined if any operand is float.

Accumulator is allowed to be the source of this instruction, but source modifier is not allowed for an accumulator source.

When accumulator is the destination of this instruction, only the low bits corresponding to the data type (16 bits for word or 32 bits for dword integer instruction) in the accumulator contain the correct results. The internal extra-precision bits as well as the sign bit of the accumulator are undefined. Consequently, there are restrictions for subsequent instructions that use the data in the accumulator register created from the previous logical instruction.

- Only logical and data move instructions are allowed to source the accumulator. Results of other instructions (e.g. arithmetic or shift) are undefined.
- When the accumulator is the source of a data move (mov or sel) instruction, the destination operand must be of integer type (e.g. no conversion to float) and this instruction cannot have saturation instruction modifier.



### 8.3.3 asr – Arithmetic Shift Right

Opcode	Instruction	Description
12 (0x0C)	asr <dst> <src0> <src1>	Performing component-wise arithmetic right shift of <src0> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[INT]	[INT]

#### Format:

[(<pred>)] asr[.<cmode>] (<exec\_size>) <dst> <src0> <src1>

#### Syntax:

[(<pred>)] asr[.<cmode>] (<exec\_size>) reg reg reg  
 [(<pred>)] asr[.<cmode>] (<exec\_size>) reg reg imm32

#### Pseudocode:

```

Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.channel[n] == 1) {
        if (src0.chan[n] >= 0) {
            dst.chan[n] = src0.chan[n] >> src1.chan[n];
        } else {
            int maskLSB = pow(2, src1.chan[n]) - 1;
            if (maskLSB & src0.chan[n] == 0) {
                dst.chan[n] = sign(src0.chan[n]) *
                    ((abs)src0.chan[n] >> src1.chan[n]);
            } else {
                dst.chan[n] = sign(src0.chan[n]) *
                    ((abs)src0.chan[n] >> src1.chan[n]) - 1;
            }
        }
    }
}

```



**Description:**

The *asr* instruction performs component-wise arithmetic right shift of <src0> and storing the results in <dst>. Arithmetic right shift performs sign-extension by repeating the MSB of each data channel of <src0>. The amount of bit shift is provided by <src1>. Only the 5 LSBs of each channel of <src1> are used as an unsigned integer value. The rest of MSBs of <src1> data channels are ignored.

Operands for this instruction can be signed or unsigned integers, but cannot be floating point type. 5-bit shifting applies to packed-dword mode and packed-word mode. For packed word mode, the accumulators have 33 bits per channel.

This instruction is effectively a power-of-2 integer divide with truncate in 2's compliment form. Truncate in 2's compliment form is also known as downward rounding – closest integer that is smaller than or equal to the result. For example, regardless of the bit shift amount in <src1>, the result of arithmetic right-shift of -1 (<src0>) is always -1.

**Restrictions:**

This instruction does not work with float type operands.

### 8.3.4 avg – Average

Opcode	Instruction	Description
66 (0x42)	avg <dst> <src0> <src1>	Component-wise averaging of <src0> and <src1> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[INT]	[INT]

**Format:**

[[<pred>]] avg[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[[<pred>]] avg[.<cmod>] (<exec\_size>) reg reg reg  
 [[<pred>]] avg[.<cmod>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = (src0.chan[n] + src1.chan[n] + 1) >> 1;
    }
}
```

**Description:**

The `avg` instruction performs component-wise integer average of `<src0>` and `<src1>` and stores the results in `<dst>`. An integer average uses integer upward rounding. It is equivalent to increment one to the addition of `<src0>` and `<src1>` and then apply an arithmetic right shift to this intermediate value.

**Restrictions:**

This instruction only applies to integer data types. The behavior is undefined if any operand is float.

### 8.3.5 `break` – Break [DevSNB]

Opcode	Instruction	Description
40 (0x28)	Break <JIP><UIP>	Terminating enabled execution channels and conditionally breaking out from the inner most loop.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•					

**Format:**

`[(<pred>)] break (<exec_size>) <JIP> <UIP>`

**Syntax:**

`[(<pred>)] break (<exec_size>) imm16 imm16`

**Pseudocode:**

```

Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.channel[n] == 1) {
        if (PMask[n] == 1) {
            PcIP[n] = IP + <UIP>;
        }
        else {
            PcIP[n] = IP+1;
        }
    }
}
if (PcIP != (IP + 1)) { //all channels
    Jump(IP + <JIP>);
}

```



**Description:**

The *break* instruction is used to early-out from the inner most loop, or early-out from the inner switch block.

When used in a loop, upon execution, the *break* instruction terminates the loop for all execution channels enabled by PMask. This is performed by updating the per channel IP to the <UIP>. In case of all the enabled channels hit the *break* instruction, a jump will be performed to the instruction based <JIP>. <JIP> should be the offset to the end of the inner most conditional or loop block, <UIP> should be the offset to the first instruction after the loop block.

When *break* used in a switch block, upon execution, the *break* instruction terminates the current enabled channels for the rest of the switch code block. In case of all the channels hit the *break* instruction, a jump will be performed to the instruction based on <JIP>. <JIP> should be the offset to the first instruction after the switch block. <UIP> should be the instruction after the entire switch block. In case there are multiple *breaks* inside a single case block, the last *break* instruction must be unconditional, and all the previous conditional *break* instructions must have the <JIP> point to the last unconditional *break* instruction. Each case block must end with a *break* instruction.

The following table describes the 2 16-bit instruction pointer offset. Both the <JIP> and <UIP> are signed 16-bit numbers, added to IP pre-increment In binary, <JIP> and <UIP> are at location <src1> and must be of type W (signed word integer).

Bit	Description
31:16	<b>UIP (Update Target Offset).</b> This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the channel. Format = S15. Signed integer in 2's compliment
15:0	<b>JIP (Jump Target Offset).</b> This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the instruction. Format = S15. Signed integer in 2's compliment

If SPF is ON, none of the PciIP is updated. The UIP must be used to update the execution IP, the JIP is not used in this case.

### 8.3.6 call – Call [DevSNB]

Opcode	Instruction	Description
44 (0x2c)	call <dst> <JIP>	Subroutine call.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•					DW

**Format:**

```
[(<pred>)] call (2) <dst> <JIP>
```

**Syntax:**

```
[(<pred>)] call (2) reg imm16
```

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < 16; n++) {
    if (WrEn.channel[n] == 1) {
        PclIP[n] = IP + <JIP>;
        CallMask[n] = 1;
    } else {
        PclIP[n] = IP + 1;
        CallMask[n] = 0;
    }
}

if (PclIP[n] != (IP+1)) { // any channel jumped
    dst.chan[0] = IP+1;
    dst.chan[1] = CallMask;
    Jump(IP + <JIP>);
}
```

**Description:**

The *call* instruction does the jump to the subroutine. It can be predicated or non-predicated. If non-predicated, all enabled channels will jump to the subroutine. If predicated, only the channels enabled by the PMask will jump to the subroutine, the rest of the channels will move to the next instruction after the call instruction, if none of the channels jump into the subroutine, the *call* instruction will be treated as a NOP.

In case of a jump, the *call* instruction stores the return IP onto the first dword of the destination register, the CallMask is stored in the second dword of the destination register.

If SPF is ON, none of the PclIP is updated.

The following table describes the 16-bit jump offset <JIP>. <JIP> must be an immediate operand, whereas it is a signed 16-bit number. When a jump occurs, this value is added to IP pre-increment. In DevSNB instruction binary, <JIP> is at location <src1> and must be of type W (signed word integer).

When SPF is on, the predication control must be scalar.



Bit	Description
31:16	<b>Reserved:</b> MBZ
15:0	<b>JIP (Jump Instruction Count).</b> This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the instruction. Format = S15. Signed integer in 2's compliment

**Restrictions:**

- The *call* instruction must have DWord source and destination type, and the destination must be QWord aligned.
- The source0 regioning control must be <2;2,1>, and the execution size must be 2.
- QtrCtrl is not allowed for this instruction.
- **[Errata][DevSNB]** When predication is required for a CALL, the instruction must be invoked inside an “if-endif” block.
- [Errata] [DevSNB-**Astep**] The call instruction does not push the correct callmask onto the stack if it is predicated or resides in a conditional code block.
- [Errata] [DevSNB-**Astep**] The destination must be Oword aligned.

### 8.3.7 case – Case [DevSNB]

Opcode	Instruction	Description
38 (0x26)	case <dst> <src0> <src1> <JIP>	Signifying the start of an case block of code.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•		•			

**Format:**

```
[(<pred>)] case (<exec_size>) null null null <JIP>
case[.<cmod>] (<exec_size>) null <src0> <src1> <JIP>
```

**Syntax:**

```
[(<pred>)] case (<exec_size>) null null null imm16
case[.<cmod>] (<exec_size>) null reg reg imm16
case[.<cmod>] (<exec_size>) null reg imm32 imm16
```

**Pseudocode:**

```
Evaluate(WrEn);
```





```
for (n = 0; n < 16; n++) {
    if (WrEn.channel[n] == 1) {
        if (<cmode> == 0) { // no embedded compare
            if (PMask.channel[n] == 0) {
                PcIP[n] = IP + <JIP>;
            }
            else
            {
                PcIP[n] = IP + 1;
            }
        }
        else { // with embedded compare
            if (cmode.channel[n] == 0) {
                PcIP[n] = IP + <JIP>;
            }
            else
            {
                PcIP[n] = IP + 1;
            }
        }
    }
}
if (<cmode> == 0) { // no embedded compare
    if (PcIP != (IP+1)) { // all channels false
        Jump(IP + <JIP>);
    }
}
else { // with embedded compare
    if (PcIP != (IP+1)) { // all channels false
        Jump(IP + <JIP>);
    }
}
```

**Description:**

The *case* instruction starts an *case/break* code block. It restricts execution within the conditional block to only those channels that were enabled via either the predicate control or the condition from <cmode>.

Each *case* instruction must have a matching *break* instruction.

If all channels are inactive (for the *case/break* block), a jump is performed of the relative distance as specified in the instruction. This jump must be to right after the matching *break* instruction when present, or otherwise to the end of *switch* code block.



The following table describes the 16-bit exit code <JIP>. <JIP> must be an immediate operand, whereas it is a signed 16-bit number. When a jump occurs, this value is added to IP pre-increment. In DevSNB instruction binary, <JIP> is at location <dst> and must be of type W (signed word integer).

Bit	Description
31:16	<b>Reserved:</b> MBZ
15:0	<b>JIP (Jump Instruction Count).</b> This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the instruction. Format = S15. Signed integer in 2's compliment

If SPF is ON, none of the PclIP is updated.

**Restrictions:**

To use embedded compares, the predicate control field for this instruction must be zero, and the conditional modifier field must be none zero. The destination must follow the rules below:

1. Must has the same element size as source0
2. Must have horizontal stride of 1

To use predicated case instruction, the conditional modifier field must be zero.

The case code block must end with a *break* instruction.

The case instruction must be used within the switch/endswitch code block

3. Accumulator can not be used as source or destination for this instruction.



### 8.3.8 cmp – Compare

Opcode	Instruction	Description
16 (0x10)	cmp.<cmode> <dst> <src0> <src1>	Component-wise comparison of <src0> and <src1> according to conditional modifier in <cmode> and storing the results in flag register in <cmode> and <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•		•	•	[FLT]	[FLT]
•		•	•	[INT]	[INT]

**Format:**

[(<pred>)] cmp[.<cmode>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] cmp[.<cmode>] (<exec\_size>) reg reg reg  
 [(<pred>)] cmp[.<cmode>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```

Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    bitMask[n] = 0;
    if (WrEn.chan[n] == 1) {
        results[n] = src0.chan[n] - src1.chan[n];
        bitMask[n] = Condition(results[n]);
        dst.chan[n][0] = bitMask[n];
    }
}
flag# = bitMask;
  
```

**Description:**

The *cmp* instruction performs component-wise comparison of <src0> and <src1> and stores the results in the selected flag register and in <dst>. It takes component-wise subtraction of <src0> and <src1>, evaluating the conditional code (excluding NS signal) based on the conditional modifier, and storing the conditional bits in bit-packed form in the destination flag register and, optionally, in vector form in the LSB of the channels in <dst>. Conditional modifier field cannot be 0000b, i.e., it must be one of the defined conditional modifier codes. Destination operand can be a GRF, an MRF or a null register. If it is not null, for the enabled channels, the LSB of the result in the destination channel contains the flag value for the



channel. The other bits are undefined. When the instruction operates on packed word format, one GRF register may store up to 16 such comparison results. In dword format, one GRF may store up to 8 results. When the register is used later as a vector of Booleans, as only LSB at each channel contains meaning data, software should make sure all higher bits are masked out (e.g. by 'and-ing' an 0x01 constant).

If <exec\_size> is equal or less than 8, when 'SecHalf' option flag is not set, only the lower 8 bits of the selected flag register is updated; otherwise, the higher 8 bits are updated.

When at least one of the source operands is float, the *cmp* instruction obeys the floating point rules detailed in the tables in the Floating Point Mode section of Data Type chapter.

**Restrictions:**

Destination operand cannot be an ARF register, including accumulator.

Saturation modifier cannot be set in this instruction.

Dword integer source is not allowed for this instruction in float execution mode. In other words, if one source is of type float (:f, :vf), the other source cannot be of type dword integer (:ud or :d).

**[DevSNB Errata]** The *cmp* instruction can not use accumulator as source if source modifier is used, all other cases are ok.

**[DevSNB Errata]** Integer compare can not use accumulator as source.

### 8.3.9 cmpn – Compare NaN

Opcode	Instruction	Description
17 (0x11)	cmpn.<cmop> <dst> <src0> <src1>	Performing component-wise special NaN comparison of <src0> and <src1> according to conditional modifier in <cmop> and storing the results in flag register in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•		•	•	[FLT]	[FLT]

**Format:**

[(<pred>)] cmpn[.<cmop>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] cmpn[.<cmop>] (<exec\_size>) reg reg reg  
 [(<pred>)] cmpn[.<cmop>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    bitMask[n] = 0;
    if (EMask.chan[n] == 1) {
        results[n] = src0.chan[n] - src1.chan[n];
        bitMask[n] = ConditionNaN(results[n]);
        dst.chan[n][0] = bitMask[n];
    }
}
flag# = bitMask;
```

**Description:**

The *cmpn* instruction performs component-wise special-NaN comparison of <src0> and <src1> and stores the results in the selected flag register and in <dst>. It takes component-wise subtraction of <src0> and <src1>, evaluating the conditional signals including NS based on the conditional modifier, and storing the conditional flag bits in bit-packed form in the destination flag register and, optionally, in vector form in the LSB of the channels in <dst>. Conditional modifier field cannot be 0000b, i.e., it must be one of the defined conditional modifier codes. Destination operand can be a GRF, an MRF or a null register. If it is not null, for the enabled channels, the LSB of the result in the destination channel contains the flag value for the channel. The other bits are undefined. When the instruction operates on packed word format, one GRF register may store up to 16 such comparison results. In dword format, one GRF may store up to 8 results. When the register is used later as a vector of Booleans, as only LSB at each channel contains meaning data, software should make sure all higher bits are masked out (e.g. by 'and-ing' an 0x01 constant).

If <exec\_size> is equal or less than 8, when 'SecHalf' option flag is not set, only the lower 8 bits of the selected flag register is updated; otherwise, the higher 8 bits are updated.

When at least one of the source operands is float, the *cmpn* instruction obeys the floating point rules detailed in the tables in the Floating Point Mode section of Data Type chapter.

This instruction is similar to *cmp*. The only difference is that if the second source operand <src1> is a NaN, the result for any conditional modifier except *.nz* is true (see Section 4.3.1.4 for details).

For integer operands, *cmpn* and *cmp* are identical.

**Restrictions:**

Destination operand cannot be an ARF register, including accumulator.

Saturation modifier cannot be set in this instruction.

Dword integer source is not allowed for this instruction in float execution mode. In other words, if one source is of type float (:f, :vf), the other source cannot be of type dword integer (:ud or :d).

**DevSNB Errata]** Integer compare can not use accumulator as source.



### 8.3.10 cont – Continue [DevSNB]

Opcode	Instruction	Description
41 (0x29)	cont <JIP><UIP>	Temporarily disabling enabled execution channels for the remainder of the inner most loop and conditionally jumping to the last instruction (while) of the loop.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•					

**Format:**

[(<pred>)] cont (<exec\_size>) <JIP> <UIP>

**Syntax:**

[(<pred>)] cont (<exec\_size>) imm16 imm16

**Pseudocode:**

```

Evaluate(WrEn);
for (n = 0; n < 16; n++) {
    if (WrEn.channel[n] == 1) {
        if (PMask[n] == 1) {
            PcIP[n] = IP + <UIP>;
        }
        else {
            PcIP[n] = IP + 1;
        }
    }
}
if (PcIP != (IP+1)) { //all channel true
    Jump(IP + <JIP>);
}

```

**Description:**

The *cont* instruction disables execution for the subset of channels for the remainder of the current loop iteration. Channels remain disabled until right before the *while* instruction or right before the condition check code block for the *while* instruction. In case of all enabled channels hit this instruction, a jump is made a distance of <JIP> where execution continues.



The following table describes the two 16-bit exit code, <JIP> and <UIP>. The field are signed 16-bit numbers, added to IP pre-increment. The <UIP> should always point to the loop's associated 'while' instruction, and the <JIP> should point to the last instruction of the inner most conditional block if the *cont* instruction is inside a conditional block. In case of the *break* instruction directly under the loop, the <JIP> and the <UIP> will be the same. In DevSNB binary, <JIP> and <UIP> are at location <src1> and must be of type W (signed eword integer).

Bit	Description
31:16	<b>UIP (Update Target Offset).</b> This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the channel. Format = S15. Signed integer in 2's compliment
15:0	<b>JIP (Jump Target Offset).</b> This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the instruction. Format = S15. Signed integer in 2's compliment

If SPF is ON, none of the PciP is updated. The UIP must be used to update the execution IP, the JIP is not used in this case.

(1) PMask here is for all the channels enabled for the *cont* instruction.

### 8.3.11 dp2 – Dot Product 2

Opcode	Instruction	Description
87 (0x57)	Dp2 <dst> <src0> <src1>	Performing two-wide dot product in four-tuples of <src0> and <src1> and storing the replicated results in four-tuples in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]

**Format:**

[[<pred>]] dp2[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[[<pred>]] dp2[.<cmod>] (<exec\_size>) reg reg reg  
[[<pred>]] dp2[.<cmod>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n+=4) {
    fTmp = src0.chan[n] * src1.chan[n]
        + src0.chan[n+1] * src1.chan[n+1];
```



```

    if (WrEn.chan[n] == 1) dst.chan[n] = fTmp;
    if (WrEn.chan[n+1] == 1) dst.chan[n+1] = fTmp;
    if (WrEn.chan[n+2] == 1) dst.chan[n+2] = fTmp;
    if (WrEn.chan[n+3] == 1) dst.chan[n+3] = fTmp;
}

```

**Description:**

The *dp2* instruction performs a two-wide dot-product on four-tuple vector basis and storing the same scalar result per four-tuple to all four channels in <dst>. This instruction is similar to *dp4* except that every third and fourth elements of <src0> (post-source-swizzle if present) are not involved in the computation.

Special care has been taken in the hardware such that if the resulting value for a given group of four channels is 0.0f, the sign of the result correctly reflects the input data of the first two channels of the group of four.

**Restrictions:**

Source operands cannot be an accumulator register.

Execution size cannot be less than 4.

This instruction does not support integer operation.

Horizontal stride must be 1.

[Pre-DevSNB] The results are NOT stored in the accumulator register. This instruction does implicitly update accumulator register, however with undefined values.

[DevSNB+] The results can be stored in the accumulator register, the channels are wirrten the same way as the <dst> register with replication.

### 8.3.12 dp3 – Dot Product 3

Opcode	Instruction	Description
86 (0x56)	dp3 <dst> <src0> <src1>	Performing three-wide dot product in four-tuples of <src0> and <src1> and storing the replicated results in four-tuples in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]



**Format:**

[(<pred>)] dp3[.<cmop>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] dp3[.<cmop>] (<exec\_size>) reg reg reg  
[(<pred>)] dp3[.<cmop>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n+=4) {
    fTmp = src0.chan[n] * src1.chan[n]
          + src0.chan[n+1] * src1.chan[n+1]
          + src0.chan[n+2] * src1.chan[n+2];
    if (WrEn.chan[n] == 1) dst.chan[n] = fTmp;
    if (WrEn.chan[n+1] == 1) dst.chan[n+1] = fTmp;
    if (WrEn.chan[n+2] == 1) dst.chan[n+2] = fTmp;
    if (WrEn.chan[n+3] == 1) dst.chan[n+3] = fTmp;
}
```

**Description:**

The *dp3* instruction performs a three-wide dot-product on four-tuple vector basis and storing the same scalar result per four-tuple to all four channels in <dst>. This instruction is similar to *dp4* except that every fourth element of <src0> (post-source-swizzle if present) is not involved in the computation.

Special care has been taken in the hardware such that if the resulting value for a given group of 4 channels is 0.0f, the sign of the result correctly reflects the input data of the first three channels of the group of four.

**Restrictions:**

Source operands cannot be an accumulator register.

Execution size cannot be less than 4.

This instruction does not support integer operation.

Horizontal stride must be 1.

[Pre-DevSNB]The results are NOT stored in the accumulator register. This instruction does implicitly update accumulator register, however with undefined values.

[DevSNB+]The results can be stored in the accumulator register, the channels are wirrtten the same way as the <dst> register with replication.



### 8.3.13 dp4 – Dot Product 4

Opcode	Instruction	Description
84 (0x54)	dp4 <dst> <src0> <src1>	Performing four-wide dot product of <src0> and <src1> and storing the four-wide replicated results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]

**Format:**

[(<pred>)] dp4[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] dp4[.<cmod>] (<exec\_size>) reg reg reg  
 [(<pred>)] dp4[.<cmod>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n+=4) {
    fTmp = src0.chan[n] * src1.chan[n]
          + src0.chan[n+1] * src1.chan[n+1]
          + src0.chan[n+2] * src1.chan[n+2]
          + src0.chan[n+3] * src1.chan[n+3];
    if (WrEn.chan[n] == 1) dst.chan[n] = fTmp;
    if (WrEn.chan[n+1] == 1) dst.chan[n+1] = fTmp;
    if (WrEn.chan[n+2] == 1) dst.chan[n+2] = fTmp;
    if (WrEn.chan[n+3] == 1) dst.chan[n+3] = fTmp;
}
```

**Description:**

The *dp4* instruction performs a four-wide dot-product on four-tuple vector basis and storing the same scalar result per four-tuple to all four channels in <dst>.

**Restrictions:**

- Source operands cannot be an accumulator register.
- Execution size cannot be less than 4.
- This instruction does not support integer operation.



Horizontal stride must be 1.

[Pre-DevSNB]The results are NOT stored in the accumulator register. This instruction does implicitly update accumulator register, however with undefined values.

[DevSNB+]The results can be stored in the accumulator register, the channels are wirrten the same way as the <dst> register with replication.

### 8.3.14 dph –Dot Product Homogeneous

Opcode	Instruction	Description
85 (0x55)	dph <dst> <src0> <src1>	Performing four-wide homogeneous dot product of <src0> and <src1> and storing the four-wide replicated results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]

**Format:**

[(<pred>)] dph[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] dph[.<cmod>] (<exec\_size>) reg reg reg  
 [(<pred>)] dph[.<cmod>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n+=4) {
    fTmp = src0.chan[n] * src1.chan[n]
          + src0.chan[n+1] * src1.chan[n+1]
          + src0.chan[n+2] * src1.chan[n+2]
          + src1.chan[n+3];
    if (WrEn.chan[n] == 1) dst.chan[n] = fTmp;
    if (WrEn.chan[n+1] == 1) dst.chan[n+1] = fTmp;
    if (WrEn.chan[n+2] == 1) dst.chan[n+2] = fTmp;
    if (WrEn.chan[n+3] == 1) dst.chan[n+3] = fTmp;
}
```



**Description:**

The *dph* instruction performs a four-wide homogeneous dot-product on four-tuple vector basis and storing the same scalar result per four-tuple to all four channels in <dst>. This instruction is similar to dp4 except that every fourth element of <src0> (post-source-swizzle if present) is forced to 1.0f.

**Restrictions:**

Source operands cannot be an accumulator register.

Execution size cannot be less than 4.

This instruction does not support integer operation.

Horizontal stride must be 1.

[Pre-DevSNB]The results are NOT stored in the accumulator register. This instruction does implicitly update accumulator register, however with undefined values.

[DevSNB+]The results can be stored in the accumulator register, the channels are wirrten the same way as the <dst> register with replication.

### 8.3.15 else – Else [DevSNB]

Opcode	Instruction	Description
36 (0x24)	else <JIP>	An optional statement within an if/else/endif block of code.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types

**Format:**

else (<exec\_size>) <JIP>

**Syntax:**

else (<exec\_size>) imm16

**Pseudocode:**

```

Evaluate(WrEn);
for (n = 0; n < 16; n++) {
    if (WrEn.channel[n] == 1) {
        PclP[n] = IP + <JIP>;
    }
}

```



```

if (PcIP != (IP+1)) { // for all channels
    Jump(IP + <JIP>);
}

```

**Description:**

The *else* instruction is an optional statement within an *if/else/endif* block of code. It restricts execution within the *else/endif* portion to the opposite set of channels enabled under the *if/else* portion. Channels which were inactive prior to entering the *if/endif* block remain inactive throughout the entire block.

All enabled channels upon arriving the *else* instruction will be redirected to the matching *endif*. If all channels are redirected (by *else* or before *else*), a relative jump is performed to the location specified by <JIP>. The jump target should be the the matching *endif* instruction for that conditional block.

The following table describes the 16-bit <JIP>. In DevSNB binary, <JIP> is at location <dst> and must be of type W (signed word integer). <JIP> must be an immediate operand, it is a signed 16-bit number and is intended to be forward referencing. This value is added to IP pre-increment.

Bit	Description
31:16	<b>Reserved:</b> MBZ
15:0	<b>JIP (Jump Instruction Count).</b> This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the instruction. Format = S15. Signed integer in 2's compliment

If SPF is ON, none of the PcIP is updated.

**Restrictions:**

Predication is not allowed.

### 8.3.16 endif – End-If

Opcode	Instruction	Description
37 (0x24)	endif <JIP>	Restoring execution to those data channels that were active prior to the <i>if/else/endif</i> block. Jump to next hop point if all channels are disabled.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types



**Format:**

endif <JIP>

**Syntax:**

endif imm16

**Pseudocode:**

```
Evaluate(WrEn);  
if (WrEn == 0) { // all channels false  
    Jump(IP + <JIP>);  
}
```

**Description:**

The *endif* instruction terminates an if/else/endif block of code. It restores the execution to these data channels that were active prior to the if/else/endif block.

The *endif* instruction is also used to hop out nested conditionals by jumping to the end of the next outer conditional block when all channels are disabled.

The following table describes the 16-bit jump target offset <JIP>. In DevSNB binary, <JIP> is at location <src1> and must be of type W (signed word integer). <JIP> must be an immediate operand.

Bit	Description
31:16	<b>Reserved:</b> MBZ
15:0	<b>JIP (Jump Target Offset).</b> This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the instruction. Format = S15. Signed integer in 2's compliment

Execution size is ignored for this instruction.

If SPF is ON, none of the PclIP is updated.

**Restrictions:**

Predication is not allowed.

Accumulator can not be used as source or destination for this instruction.



### 8.3.17 fork – Fork [DevSNB]

Opcode	Instruction	Description
46 (0x2E)	fork <src0> <JIP>	Set active set of 16 channels according to <src0>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•					

**Format:**

[(<pred>)] fork (<exec\_size>) <src0> <JIP>

**Syntax:**

[(<pred>)] fork (<exec\_size>) reg imm16

**Pseudocode:**

```
if (src0 == NULL) {
    SPF = 1;
} else {
    SPF = 0;

    for (n = 0; n < exec_size; n++) {
        if (src0 == DMask) {
            if (DMask.channel[n] == 0) {
                PciIP[n] = IP + <JIP>;
            }
            else {
                PciIP[n] = IP+1;
            }
        }
        if (src0 == VMask) {
            if (VMask.channel[n] == 0) {
                PciIP[n] = IP + <JIP>;
            }
            else {
                PciIP[n] = IP+1;
            }
        }
    }
}
```



```
    }  
  
    if (PcIP != (IP+1)) { // for all channels  
        Jump(IP + <JIP>);  
    }  
}
```

### Description:

The *fork* instruction starts a *fork* block of code. It restricts execution within the *fork* block to only those channels that were enabled via select register (dispatch mask or vector mask) and further qualified by the predicate control.

If all channels are inactive (for the *fork* block), a jump is performed of the relative distance as specified in the instruction.

The following table describes the 16-bit exit code <JIP>. <JIP> must be an immediate operand, whereas JumpIP is a signed 16-bit number. When a jump occurs, this value is added to IP pre-increment. In DevSNB instruction binary, <JIP> is at location <dst> and must be of type W (signed word integer).

If the register specified is NULL, fork simply set SPF = 1.

The *fork* instruction should be used to turn SPF off in a middle of a program, it will bring the PcIPs the the current ExIP for continuing execution.

Bit	Description
15:0	<b>JIP (Jump Instruction Count)</b> . This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the instruction. Format = S15. Signed integer in 2's compliment

### Restrictions:

- Instruction is only allowed at global scope.
- Src0 is limited to DMask.0, Dmask.1, Vmask.0, Vmask.1 and NULL.
- The *fork* instruction should only be used to enable 32 channel execution where per-channel branch is required. The SPF control must be set during thread dispatch time based on 32 channel dispatch. Within the shader, the fork instruction should only be used to switch between the 3 modes below:
  - 32 channel execution in SPF on mode, where no per-channel branch can be performed
  - lo-16 channel execution in SPF off mode, where per-channel branch on the low 16 channels can be performed, **QtrCtrl {H1} must be used in this case.**





- hi-16 channel execution in SPF off mode, where per-channel branch on the high 16 channels can be performed, **QtrCtrl {H2} must be used in this case.**
- Qtr control is not allowed for the *fork* instruction.

### 8.3.18 frc – Fraction

Opcode	Instruction	Description
67 (0x43)	frc <dst> <src0>	Taking component-wise truncate-to-minus-infinity fraction operation of <src0> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•		•	•	[FLT]	[FLT]
•		•	•	[FLT]	[INT]

**Format:**

[(<pred>)] frc[.<cmod>] (<exec\_size>) <dst> <src0>

**Syntax:**

[(<pred>)] frc[.<cmod>] (<exec\_size>) reg reg  
 [(<pred>)] frc[.<cmod>] (<exec\_size>) reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = src0.chan[n] – floor(src0.chan[n]);
    }
}
```

**Description:**

The *frc* instruction computes, component-wise, the truncate-to-minus-infinity fractional values of <src0> and stores the results in <dst>. The results, in the range of [0.0, 1.0], are the fractional portion of the source data.

Source operand for this instruction must be of floating point type. This instruction can only operate on normalized floating source and therefore can not take accumulator as source or destination operand.

This instruction only applies to floating point operands.



Floating point fraction computation follows rules in Table 8-3 (or Table 8-4), if the current floating point mode is IEEE mode (or ALT mode).

**Table 8-3. Floating point fraction computation in IEEE mode**

<src0>	-inf	-finite	-denorm	-0	+0	+denorm	+finite	+inf	NaN
<dst>	NaN	*	+0	+0	+0	+0	*	NaN	NaN
Notes:									
* Result is in the range of [+0, 1) – not including 1.									

**Table 8-4. Floating point fraction computation in ALT mode**

<src0>	- fmax	-finite	-denorm	-0	+0	+denorm	+finite	+ fmax	**
<dst>	+0	*	+0	+0	+0	+0	*	+0	
Notes:									
* Result is in the range of [+0, 1) – not including 1.									
** Result is undefined if <src0> is {-inf, +inf, NaN}.									

**Restrictions:**

Saturation modifier does not apply to this instruction.

This instruction can not take accumulator as source or destination operand as it can only operate on normalized floating source.

This instruction does implicitly update accumulator register when enabled, however with undefined values.



### 8.3.19 halt – Halt [DevSNB]

Opcode	Instruction	Description
42 (0x2A)	halt <JIP> <UIP>	Temporarily suspending execution for all enabled execution channels.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•					

**Format:**

[[<pred>]] halt (<exec\_size>) <JIP> <UIP>

**Syntax:**

[[<pred>]] halt (<exec\_size>) imm16 imm16

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < 16; n++) {
    if (WrEn.channel[n] == 1) {
        if (PMask.channel[n] == 1) {
            PcIP[n] = IP + <UIP>;
        } else {
            PcIP[n] = (IP+1);
        }
    }
}
if (PcIP != (IP+1)) { // for all channels
    Jump(IP + <JIP>);
}
```

**Description:**

The *halt* instruction temporarily suspends execution for all enabled compute channels. Upon execution, the enabled channels will be sent to the instruction at (IP+<UIP>), if all channels are enabled at HALT, a jump to (IP+<JIP>) will be performed.

In case of the *halt* instruction not inside any conditional code block, the value of <JIP> and <UIP> should be the same. In case of the *halt* instruction inside conditional code block, the <UIP> should be the end of the program, and the <JIP> should be end of the most inner conditional code block.



The following table describes the 16-bit jump offsets. In DevSNB binary, <JIP> and <UIP> are at location <src1> and must be of type W (signed word integer). The <JIP> and <UIP> are added to IP pre-increment.

Bit	Description
31:16	<b>UIP (Update Target Offset).</b> This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the channel. Format = S15. Signed integer in 2's compliment
15:0	<b>JIP (Jump Target Offset).</b> This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the instruction. Format = S15. Signed integer in 2's compliment

If SPF is ON, none of the PciIP is updated. The UIP must be used to update the execution IP, the JIP is not used in this case.

**Restrictions:**

<dst> and <src0> must be NULL.

### 8.3.20 if – If [DevSNB]

Opcode	Instruction	Description
34 (0x22)	if <dst> <src0> <src1> <JIP>	Signifying the start of an if/else/endif block of code.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•		•			

**Format:**

```
[(<pred>)] if (<exec_size>) null null null <JIP>
if[.<cmod>] (<exec_size>) null <src0> <src1> <JIP>
```

**Syntax:**

```
[(<pred>)] if (<exec_size>) null null null imm16
if[.<cmod>] (<exec_size>) null reg reg imm16
if[.<cmod>] (<exec_size>) null reg imm32 imm16
```

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < 16; n++) {
    if (WrEn.channel[n] == 1) {
```



```
        if (<cmode> == 0) { // no embedded compare
            if (PMask.channel[n] == 0) {
                PcIP[n] = IP + <JIP>;
            }
            else {
                PcIP[n] = IP+1;
            }
        }
    else { // with embedded compare
        if (cmode.channel[n] == 0) {
            PcIP[n] = IP + <JIP>;
        }
        else {
            PcIP[n] = (IP+1);
        }
    }
}
}
}
if (<cmode> == 0) { // no embedded compare
    if (PcIP != (IP+1)) { // for all channels
        Jump(IP + <JIP>);
    }
}
else { // with embedded compare
    if (PcIP != (IP+1)) { // for all channels
        Jump(IP + <JIP>);
    }
}
```

#### **Description:**

The *if* instruction starts an *if/endif* or an *if/else/endif* block of code. It restricts execution within the conditional block to only those channels that were enabled via either the predicate control or the condition from *<cmode>*.

Each *if* instruction must have a matching *endif* instruction and may have up to one matching *else* instruction before *endif*.

If all channels are inactive (for the *if/endif* or *if/else* block), a jump is performed of the relative distance as specified in the instruction. This jump must be to right after the matching *else* instruction when present, or otherwise to the matching *endif* instruction of that conditional block.



The following table describes the 16-bit exit code <JIP>. <JIP> must be an immediate operand, whereas JumpIP is a signed 16-bit number. When a jump occurs, this value is added to IP pre-increment. In DevSNB instruction binary, <JIP> is at location <dst> and must be of type W (signed word integer).

Bit	Description
15:0	<b>JIP (Jump Instruction Count)</b> . This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the instruction. Format = S15. Signed integer in 2's compliment

If SPF is ON, none of the PciP is updated.

**Restrictions:**

To use embedded compares, the predicate control field for this instruction must be zero, and the conditional modifier field must be none zero. The destination must follow the rules below:

3. Must has the same element size as source0
4. Must have horizontal stride of 1

To use predicated *if* instruction, the conditional modifier field must be zero.

### 8.3.21 jmp\_i – Jump Indexed

Opcode	Instruction	Description
32 (0x20)	jmp_i <index>	Redirecting program execution to <index> instructions forward of the current post-incremented instruction pointer.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•					

**Format:**

[(<pred>)] jmp\_i (1) <exitcode> {NoMask}

**Syntax:**

[(<pred>)] jmp\_i (1) reg32 {NoMask}  
 [(<pred>)] jmp\_i (1) imm32 {NoMask}

**Pseudocode:**

```
Evaluate(WrEn);
if (WrEn != 0) {
    Jump(<exitcode.index> + 1); // +1 if compacted, +2 if not.
}
```



### Description:

The *jmp* instruction redirects program execution to <exitcode.index> instructions forward of the current **post-incremented** instruction pointer. <exitcode.index> is treated a signed integer value, with positive integers or zero generating forward jumps, and negative integers generating backward jumps. An <exitcode.index> value of 0 means execution continues at the instruction immediately following the *jmp* instruction, while an index value of -1 would imply an infinite loop.

Bit	Description
31:0	<b>index (Jump Index)</b> This field specifies byte aligned jump distance if a jump is taken for the instruction. Format = S31. Signed integer in 2's compliment

<exitcode> may be a scalar register or an immediate. The data type of <exitcode > must be D (signed doubleword integer).

This instruction executes regardless of the calculated WrEn at the time of issue. – To reduce hardware complexity, instruction optional control {NoMask} must be set for this instruction. This instruction invokes a thread switch after issue to allow any masks and/or IP to be resolved if necessary.

Execution size must be 1.

Predication is allowed to provide conditional jump with a scalar condition. As the execution size is 1, the first channel of PMASK (flags post prediction control and negate) is used to determine whether the jump is taken or not. If the condition is false, the jump is not taken and the IP immediately following will be executed next.

In DevSNB binary, <exitcode.index> is at location <src1>. IP register must be put (for example, by the assembler) at <dst> and <src0> locations.

If SPF is ON, none of the PcIP is updated



### 8.3.22 line – Line

Opcode	Instruction	Description
89 (0x59)	Line <dst> <src0> <src1>	Computing a component-wise line equation ( $v = p*u+q$ ) of <src0> and <src1> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]
•	•	•	•	[INT]	[INT]

**Format:**

[[<pred>]] line[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[[<pred>]] line[.<cmod>] (<exec\_size>) reg reg reg  
 [[<pred>]] line[.<cmod>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```

Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
  for (n = 0; n < exec_size; n++) {
    dwP = src0.RegNum.SubRegNum[bits4:2] // a DW aligned scalar
    dwQ = src0.RegNum.(SubRegNum[bit4])0x8 // 4-th component
    if (WrEn.chan[n] == 1) {
      dst.chan[n] = dwP * src1.chan[n] + dwQ
    }
  }
}

```

**Description:**

The *line* instruction computes a component-wise line equation ( $v = p*u+q$  where  $u/v$  are vectors and  $p/q$  are scalars) of <src0> and <src1> and storing the results in <dst>. <src1> is the input vector  $u$ . <src0> provides input scalars  $p$  and  $q$ , where  $p$  is the scalar value based on the region description of <src0> and  $q$  is the scalar value implied from <src0> region. Specifically,  $q$  is the fourth component of the 4-tuple (128-bit aligned) that  $p$  belongs to.

**Restrictions:**

This is a specialized instruction that only support execution size of 8 or 16.

<src0> region must be a replicated scalar (with HorzStride = VertStride = 0).





Dword integer source is not allowed for this instruction in float execution mode. In other words, if one source is of type float (:f, :vf), the other source cannot be of type dword integer (:ud or :d).

Source operands cannot be an accumulator register.

<src0> for *lzd* instruction has to have .0 or .4 as the subregister number.

### 8.3.23 lzd – Leading Zero Detection

Opcode	Instruction	Description
74 (0x4A)	lzd <dst> <src0>	Performing component-wise leading zero detection of <src0> and storing the results in <dst>.

Instr Comp	Imp'd Accu	Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	•		[INT]	[INT]

#### Format:

`[(<pred>)] lzd[.<cmod>] (<exec_size>) <dst> <src0>`

#### Syntax:

`[(<pred>)] lzd[.<cmod>] (<exec_size>) reg reg`

`[(<pred>)] lzd[.<cmod>] (<exec_size>) reg reg`

#### Pseudocode:

```

Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        UD udScalar = src0.chan[n];
        UD cnt = 0;
        while ( (udScalar & (1<<31)) == 0 && cnt != 32) {
            cnt ++;
            udScalar = udScalar << 1;
        }
        dst.chan[n] = cnt;
    }
}

```

#### Description:

The *lzd* instruction counts component-wise the leading zeros from <src0> and storing the resulting counts in <dst>.



This instruction only work on unsigned dword source. Source operand may be a signed or unsigned. If it is a signed integer, source modifier (abs) must be used to convert the source into an unsigned integer type.

The destination operand must also be of unsigned dword type.

**Restrictions:**

The destination operand cannot be the accumulator.

This instruction does implicitly update accumulator register when enabled, however with undefined values.

### 8.3.24 *lrp* – Linear Interpolation

Opcode	Instruction	Description
92 (0x5c)	<i>lrp</i> <dst> <src0> <src1>	Computing a component-wise <i>lrp</i> equation ( $w = u*x+v*(1-x)$ ) of vectors (x, u, v) from <src0>, <src1>, and <src2> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•		•	•	[FLT]	[FLT]

**Format:**

[(<pred>)] *lrp* [<cmod>] (<exec\_size>) <dst> <src0> <src1> <src2>

**Syntax:**

[(<pred>)] *lrp* [<cmod>] (<exec\_size>) reg reg reg

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = src1.chan[n] * src0.chan[n] + src2.chan[n] * (1 -
src0.chan[n])
    }
}
```

**Description:**

The *lrp* instruction takes component-wise multiplication of <src0> and <src1>, and adds the result to the component-wise multiplication of <src2> and (1 - <src0>), and then stores the final results in <dst>.

The *lrp* instruction uses the 3-source instruction format.

**Restrictions:**

This instruction only supports float source and destination.

Immediate source is not allowed for *Irp*.

The vertical stride is overloaded to 4 in HW for 3-src instructions.

The overflow conditional modifier is not allowed.

### 8.3.25 mac – Multiply Accumulate

Opcode	Instruction	Description
72 (0x48)	mac <dst> <src0> <src1>	Performing component-wise multiply accumulate of <src0> and <src1> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]
•	•	•	•	[INT]	[INT]

**Format:**

[(<pred>)] mac[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] mac[.<cmod>] (<exec\_size>) reg reg reg  
[(<pred>)] mac[.<cmod>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = src0.chan[n] * src1.chan[n] + acc0.chan[n]
    }
}
```

**Description:**

The *mac* instruction takes component-wise multiplication of <src0> and <src1>, adds the results with the corresponding accumulator values, and then stores the final results in <dst>.

**Restrictions:**

Accumulator is an implied source to the addition portion of the computation. Explicit source operands cannot be accumulator.



This instruction doesn't support dword integers (D or UD).

### 8.3.26 mach – Multiply Accumulate High

Opcode	Instruction	Description
73 (0x49)	mach <dst> <src0> <src1>	Performing component-wise multiply accumulation of <src0>, <src1> and accumulator register, and returning the high dword of results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
<input type="checkbox"/>			<input type="checkbox"/>	[INT]	[INT]

Format:

[(<pred>)] mach[.<cmode>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] mach[.<cmode>] (<exec\_size>) reg reg reg

[(<pred>)] mach[.<cmode>] (<exec\_size>) reg reg imm32

Pseudocode:

Evaluate(WrEn);

```
for (n = 0; n < exec_size; n++) {
```

```
  acc0.chan[n][63:0]
```

```
  = (src0.chan[n][31:16] * src1.chan[n][31:0]) << 16
```

```
  + acc0.chan[n][63:0];
```

```
  if (WrEn.channel[n] == 1) {
```

```
    dst.channel[n][31:0] = acc0.chan[n][63:32]
```

```
  }
```

```
}
```

Description:



The mach instruction performs dword integer multiply-accumulate operation and outputs the high dword (bits [63:32]). On a component by component basis, this instruction multiplies dwords in <src1> with the high words of dwords in <src0>, left-shifts the results by 16 bits, adds them with the corresponding accumulator values, and keeps the whole 64-bit results in the accumulator. It then stores the high dword (bits [63:32]) of the results in <dst>.

This instruction is intended to be used to emulate 32-bit dword integer multiplication by utilizing the large number of bits available in the accumulator. For example, the following four instructions perform vector multiplication of two 32-bit signed integer source from r2 and r3 and store the resulting vectors with high 32-bit in r5 and low 32-bit in r6.

```
mul (8) acc0:d r2.0<8;8,1>d r3.0<8;8,1>d
```

```
mach (8) rTemp<1>:d r2.0<8;8,1>d r3.0<8;8,1>d
```

```
mov (8) r5.0<1>:d rTemp:d // hi-32bits
```

```
mov (8) r6.0<1>:d acc0:d // lo-32bits
```

The MUL and MACH instruction should have all channels enabled. The first MOV should have channel enable from the destHI of IMUL, the second MOV should have the channel enable from the destLO of IMUL.

As mach is used to generate part of 64-bit dword integer results, saturation modifier should not be used. In fact, saturation modifier should not be used for any of these four instructions.

Source and destination operands must be dword integers. Source and destination must be of the same type, signed integer or unsigned integer.

If <dst> is UD, <src0> and <src1> may be UD and/or D. However, if any of <src0> and <src1> is D, source modifier, (abs), must be present to convert it to match with <dst>.

If <dst> is D, <src0> and <src1> must also be D. They cannot be UD as it may cause unexpected overflow because the computed results are limited to 64 bits.

Restrictions:

Accumulator is an implied source to the addition portion of the computation. Therefore, explicit source operands cannot be accumulator.

The AccWrEn must be set for this instruction.

The AccWrCtrl control must be set for this instruction.



### 8.3.27 mad – Multiply Add

Opcode	Instruction	Description
91 (0x5B)	mad <dst> <src0> <src1> <src2>	Component-wise multiply add of <src0>, <src1>, and <src2> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]

**Format:**

[(<pred>)] mad[.<cmod>] (<exec\_size>) <dst> <src0> <src1> <src2>

**Syntax:**

[(<pred>)] mad[.<cmod>] (<exec\_size>) reg reg reg reg

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = src1.chan[n] * src2.chan[n] + src0.chan[n]
    }
}
```

**Description:**

The *mad* instruction follows the 3-src instruction format.

The *mad* instruction takes component-wise multiplication of <src1> and <src2>, adds the results with the corresponding <src0> values, and then stores the final results in <dst>.

**Restrictions:**

This instruction only supports float source and destination.

Immediate source is not allowed for *mad*.

The vertical stride is overloaded to 4 in HW for 3-src instructions.



### 8.3.28 math – Extended Math Function [DevSNB]

Opcode	Instruction	Description
56 (0x38)	math <dst> <src0> <src1> <FC>	Component-wise extended math function on <src0> (and <src1>) to <dst> based on the <FC>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•			[FLT] [INT]	[FLT] [INT]

**Format:**

[(<pred>)] math (<exec\_size>) <dst> <src0> <src1> <FC>

**Syntax:**

[(<pred>)] math (<exec\_size>) reg reg reg imm6

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.channel[n] == 1) {
        switch FC[3:0] {
            case 1h:
                dst.channel[n] = rcp(src0.channel[n]);
            case 2h:
                dst.channel[n] = log(src0.channel[n]);
            case 3h:
                dst.channel[n] = exp(src0.channel[n]);
            case 4h:
                dst.channel[n] = sqrt(src0.channel[n]);
            case 5h:
                dst.channel[n] = rsq(src0.channel[n]);
            case 6h:
                dst.channel[n] = sin(src0.channel[n]);
            case 7h:
                dst.channel[n] = cos(src0.channel[n]);
            case 9h:
                // src0 / src1
```



```

        dst.channel[n] = fdiv(src0.channel[n], src1.channel[n]);
case Ah:
    dst.channel[n] = pow(src0.channel[n], src1/channel[n]);
case Bh:
    // src0 / src1
    idiv(src0.channel[n], src1.channel[n]);
    dst.channel[n] = quotient;
    dst+1.channel[n] = remainder;
case Ch:
    idiv(src0.channel[n], src1.channel[n]);
    dst.channel[n] = quotient;
case Dh:
    idiv(src0.channel[n], src1.channel[n]);
    dst.channel[n] = remainder;
    }
endswitch;
}
}

```

**Description:**

The *math* instruction performs extended math function on the components in <src0>, or <src0> and <src1>, and write the output to the channels of <dst>. The type of extended math function are based on the FC[3:0] encoding in the table below.

Function Control[5:0]	Description
[5]	Reserved
[4]	Reserved
[3:0]	Function Description: 0h: Reserved 1h: INV (reciprocal) 2h: LOG 3h: EXP 4h: SQRT 5h: RSQ





Function Control[5:0]	Description
	6h: SIN
	7h: COS
	8h: Reserved
	9h: FDIV
	Ah: POW
	Bh: INT DIV – return quotient and remainder
	Ch: INT DIV – return quotient only
	Dh: INT DIV – return remainder
	Eh: Reserved
	Fh: Reserved

**Restrictions:**

The *math* instruction can only use GRF registers as source(s) and destination.

The supported regioning modes for *math* instruction are align16, align1 with the following restrictions:

- Scalar source is supported.
- Source and destination horizontal stride must be 1.
- Width must be the same as execution size.
- Source and destination offset must be the same, except the case of scalar source.

The *math* instruction does not support indirect addressing modes.

For INTDIV with both quotient and remainder, the second destination register is implied from the first destination register with “+1” to get the register number.

DW and UD is the only source format supported for INT DIV, float32 is the only source format supported for all the other functions.

Mixed DW and UD sources are not allowed for the INT DIV function.

For single source math function, <src1> must be programmed as ARF-NULL register.

The *math* instruction can not be SIMD16.

The FDIV function is not supported in ALT\_MODE.



The only supported rounding mode for *math* instruction is Round to Nearest Even.

### 8.3.29 mov – Move

Opcode	Instruction	Description
1 (0x01)	mov <dst> <src0>	Component-wise move from <src0> to <dst>.

Instr Comp	Imp'd Accu	Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•		•	•	•	•	[FLT] [INT]	[FLT] [INT]

**Format:**

[[<pred>]] mov[.<cmode>] (<exec\_size>) <dst> <src0>

**Syntax:**

[[<pred>]] mov[.<cmode>] (<exec\_size>) reg reg  
 [[<pred>]] mov[.<cmode>] (<exec\_size>) reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = src0.chan[n]
    }
}
```

**Description:**

The *mov* instruction moves the components in <src0> into the channels of <dst>. If <src0> and <dst> are of different types, format conversion is performed. If <src0> is a scalar immediate, the immediate value is loaded into all channels of <dst>.

**Restrictions:**

The accumulator can be either <src0> or <dst> but not both.

This instruction does not implicitly update accumulator register.

Conditional Modifier is not allowed for this instruction.

Raw MOV is not supported for floats with ALT mode if INF/NAN is expected.



A MOV operation with NO source or destination modifiers (srcmod, saturate) or datatype conversion is considered a Raw Mov.

### 8.3.30 movi – Move Indexed

Opcode	Instruction	Description
3 (0x3)	movi <dst> <src0>	Fast component-wise indexed move from <src0> to <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT] [INT]	[FLT] [INT]

**Format:**

[(<pred>)] movi[.<cmod>] (<exec\_size>) <dst> <src0>

**Syntax:**

[(<pred>)] movi[.<cmod>] (<exec\_size>) reg reg

**Pseudocode:**

```

Evaluate(WrEn)
srcregfile = regfile(src0)
srcreg = reg(address[0])
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        srcsubreg = subreg(address[n] + addr_imm)
        dst.chan[n] = srcregfile.srcreg.srcsubreg
    }
}

```

**Description:**

The *movi* instruction performs a fast component-wise indexed move for subfields from <src0> to <dst>. The source operand must be an indirectly-addressed register. All channels of the source operand share the same register number, which is provided by the register field of the first address subregister, with a possible immediate address offset. The register fields of the subsequent address subregisters are ignored by hardware. The subregister number of a source channel is provided by the subregister field of the corresponding address subregister.

Destination register may be either a directly-addressed or an indirectly-addressed register.



This instruction effectively performs a subfield shuffling from one register to another. Up to eight subfields can be selected by an instruction.

**Restrictions:**

Source operand must be a GRF register.

Source and destination must be the same type.

This instruction does not implicitly update accumulator register.

Execution size must be less than 16 (as there are only 8 address registers)

Address register for source must be aligned to the base (a0.0).

Destination register (directly or indirectly addressed) must be half GRF aligned (i.e. 16-byte aligned).

Destination stride in unit of byte must equal to the source element size in unit of byte.

Align16 access mode is not allowed for MOVI.

All the index registers used in MOVI instruction must all point to the same GRF register.

MOVI must use 1x1 indirect regioning.

MOVI is always based on register offset zero no matter what the destination offset is. The destination offset is used in HW only to create channel enables. The first index register (a0.0) will always be used to select the first element of the destination start from offset zero. Each index register will be used to select 1 element if type is byte, each index register will be used to select 2 elements if type is word, and each index register will be used to select 4 elements if type is dword

Conditional Modifier is not allowed for this instruction.

**HW Implementation Details:**

The destination offset of the *movi* instruction is only used in HW to generate destination write enables. Each element of the destination is directly mapped to the index registers for the *movi* instruction.

For Byte *movi*, byte0 of the destination is selected by (a0.0'[4:0]), byte1 is selected by (a0.1'[4:0]), ..., and byte7 is selected by (a0.7'[4:0]). The rest of the bytes are undefined.

For Word *movi*, byte0 of the destination is selected by (a0.0'[4:1] & "0"), byte1 is selected by (a0.0'[4:1] & "1"), byte2 is selected by (a0.1'[4:1] & "0"), byte3 is selected by (a0.1'[4:1] & "1"), ..., and byte15 is selected by (a0.7'[4:1] & "1"). The rest of the bytes are undefined.

For DWord or float *movi*, byte0 of the destination is selected by (a0.0'[4:2] & "00"), byte1 is selected by (a0.0'[4:2] & "01"), byte2 is selected by (a0.0'[4:2] & "10"), byte3 is selected by (a0.0'[4:2] & "11"), byte4 is selected by (a0.1'[4:2] & "00"), byte5 is selected by (a0.1'[4:2] & "01"), ..., byte31 is selected by (a0.7'[4:2] & "11").

For all 3 conditions above,  $a0.n'[4:0] = a0.n[4:0] + \text{addr\_imm}[4:0]$ .



### 8.3.31 mul – Multiply

Opcode	Instruction	Description
65 (0x41)	mul <dst> <src0> <src1>	Performing component-wise multiplication of <src0> and <src1> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT] [INT]	[FLT] [INT]

**Format:**

[(<pred>)] mul[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] mul[.<cmod>] (<exec\_size>) reg reg reg  
[(<pred>)] mul[.<cmod>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = src0.chan[n] * src1.chan[n];
    }
}
```

**Description:**

The *mul* instruction performs component-wise multiplication of <src0> and <src1> and stores the results in <dst>.

[DevSNB]: When both <src0> and <src1> are of type D or UD, only the lower 16 bits of each element of <src0> are used. Accumulator maintains full 48-bit precision. The macro described in *mach* instruction should be used to obtain the full precision 64 bits multiplication results.

Multiplication of two floating point numbers follows rules in Table 8-5 (or



Table 8-6), if the current floating point mode is IEEE mode (or ALT mode).

**Table 8-5. Floating point multiplication of A (column) and B (row) in IEEE mode**

	-inf	-finite	-1.0	-denorm	-0	+0	+denorm	+1.0	+finite	+inf	NaN
-inf	+inf	+inf	+inf	NaN	NaN	NaN	NaN	-inf	-inf	-inf	NaN
-finite	+inf	*	-A	+0	+0	-0	-0	A	**	-inf	NaN
-1.0	+inf	-B	+1.0	+0	+0	-0	-0	-1.0	-B	-inf	NaN
-denorm	NaN	+0	+0	+0	+0	-0	-0	-0	-0	NaN	NaN
-0	NaN	+0	+0	+0	+0	-0	-0	-0	-0	NaN	NaN
+0	NaN	-0	-0	-0	-0	+0	+0	+0	+0	NaN	NaN
+denorm	NaN	-0	-0	-0	-0	+0	+0	+0	+0	NaN	NaN
+1.0	-inf	B	-1.0	-0	-0	+0	+0	+1.0	B	+inf	NaN
+finite	-inf	**	-A	-0	-0	+0	+0	A	*	+inf	NaN
+inf	-inf	-inf	-inf	NaN	NaN	NaN	NaN	+inf	+inf	+inf	NaN
NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN

Note:

\* Result may be {+finite, +inf (overflow)}

\*\* Result may be {-inf (overflow), -finite}



**Table 8-6. Floating point multiplication of A (column) and B (row) in ALT mode**

	- fmax	-finite	-1.0	-denorm	-0	+0	+denorm	+1.0	+finite	+fmax	***
- fmax	+fmax	+fmax	+fma x	-0	-0	+0	+0	-fmax	-fmax	-fmax	
-finite	+fmax	*	-A	+0	+0	-0	-0	A	**	-fmax	
-1.0	+fmax	-B	+1.0	+0	+0	-0	-0	-1.0	-B	-fmax	
-denorm	+0	+0	+0	+0	+0	-0	-0	-0	-0	-0	
-0	+0	+0	+0	+0	+0	-0	-0	-0	-0	-0	
+0	-0	-0	-0	-0	-0	+0	+0	+0	+0	+0	
+denorm	-0	-0	-0	-0	-0	+0	+0	+0	+0	+0	
+1.0	-fmax	B	-1.0	-0	-0	+0	+0	+1.0	B	+fmax	
+finite	-fmax	**	-A	-0	-0	+0	+0	A	*	+fmax	
+fmax	-fmax	-fmax	- fmax	-0	-0	+0	+0	+fmax	+fmax	+fmax	
***											

Note:

- \* Result may be {+finite, +fmax (overflow)}
- \*\* Result may be {-fmax (overflow), -finite}
- \*\*\* Result is undefined If any of A and/or is {-inf, +inf, NaN}

**Restrictions:**

Source operands cannot be an accumulator register.

When operating on integers with at least one of the source being a dword type (signed or unsigned), the destination cannot be a float (implementation note: the data converter only looks at the lower 34 bits of the results).

Dword integer source is not allowed for this instruction in float execution mode. In other words, if one source is of type float (:f, :vf), the other source cannot be of type dword integer (:ud or :d).

When operating on integers with at least one of the source being a dword type (signed or unsigned), the Overflow and Sign flags are undefined. Therefore, conditional modifier and instruction operation '.sat' cannot be used.

[DevSNB]: When multiple a DW and a W, the W has to be on src0, and the DW has to be on src1.



### 8.3.32 nop – No Operation

Opcode	Instruction	Description
126 (0x7E)	nop	Issuing an dummy instruction and performing no operation.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types

**Format:**

nop

**Syntax:**

nop

**Pseudocode:**

n/a

**Description:**

The *nop* instruction takes an instruction dispatch but performs no operation. It may be used for assembly patching in memory, or be used to insert an instruction delay in the program sequence.

The *nop* instruction takes no operands, no instruction modifier, no conditional modifier and no predication.





### 8.3.33 not – Logic Not

Opcode	Instruction	Description
4 (0x04)	not <dst> <src0>	Performing component-wise logic NOT of <src0> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•		•	•	[INT]	[INT]

#### Format:

[(<pred>)] not[.<cmode>] (<exec\_size>) <dst> <src0>

#### Syntax:

[(<pred>)] not[.<cmode>] (<exec\_size>) reg reg  
[(<pred>)] not[.<cmode>] (<exec\_size>) reg imm32

#### Pseudocode:

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = !src0.chan[n]
    }
}
```

#### Description:

The *not* instruction performs logical NOT operation (or one's complement) of <src0> and storing the results in <dst>.

Source modifiers are allowed.

Accumulator register is allowed to be the destination of this instruction with the restrictions listed below.

#### Restrictions:

This instruction does not work with float type operands.

Sign (SN) and Overflow (OF) conditions are undefined for this logic instruction. Consequently, saturation modifier (.sat) is not allowed.

This instruction does not implicitly update accumulator register.

Accumulator is allowed to be the source of this instruction, but source modifier is not allowed for an accumulator source.



When accumulator is the destination of this instruction, only the low bits corresponding to the data type (16 bits for word or 32 bits for dword integer instruction) in the accumulator contain the correct results. The internal extra-precision bits as well as the sign bit of the accumulator are undefined. Consequently, there are restrictions for subsequent instructions that use the data in the accumulator register created from the previous logical instruction.

- Only logical and data move instructions are allowed to source the accumulator. Results of other instructions (e.g. arithmetic or shift) are undefined.
- When the accumulator is the source of a data move (mov or sel) instruction, the destination operand must be of integer type (e.g. no conversion to float) and this instruction cannot have saturation instruction modifier.

### 8.3.34 or – Logic Or

Opcode	Instruction	Description
6 (0x06)	or <dst> <src0> <src1>	Performing component-wise logic OR of <src0> and <src1> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•		•	•	[INT]	[INT]

#### Format:

[(<pred>)] or[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

#### Syntax:

[(<pred>)] or[.<cmod>] (<exec\_size>) reg reg reg  
 [(<pred>)] or[.<cmod>] (<exec\_size>) reg reg imm32

#### Pseudocode:

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = src0.chan[n] | src1.chan[n];
    }
}
```

#### Description:

The *or* instruction performs component-wise logic OR operation between <src0> and <src1> and stores the results in <dst>.

Source modifiers are allowed.



Accumulator register is allowed to be the destination of this instruction with the restrictions listed below.

**Restrictions:**

Sign (SN) and Overflow (OF) conditions are undefined for this logic instruction. Consequently, saturation modifier (.sat) is not allowed.

This instruction does not work with float type operands.

Accumulator is allowed to be the source of this instruction, but source modifier is not allowed for an accumulator source.

This instruction does not implicitly update accumulator register.

When accumulator is the destination of this instruction, only the low bits corresponding to the data type (16 bits for word or 32 bits for dword integer instruction) in the accumulator contain the correct results. The internal extra-precision bits as well as the sign bit of the accumulator are undefined. Consequently, there are restrictions for subsequent instructions that use the data in the accumulator register created from the previous logical instruction.

- Only logical and data move instructions are allowed to source the accumulator. Results of other instructions (e.g. arithmetic or shift) are undefined.
- When the accumulator is the source of a data move (mov or sel) instruction, the destination operand must be of integer type (e.g. no conversion to float) and this instruction cannot have saturation instruction modifier.

### 8.3.35 pln – Plane

Opcode	Instruction	Description
90 (0x5A)	PLN <dst> <src0> <src1>	Computing a component-wise plane equation ( $w = p*u+q*v+r$ ) of scalar (p, q, r) from <src0> and vector <src1> (and implied vector <src2>) and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]

**Format:**

[(<pred>)] pln[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] pln[.<cmod>] (<exec\_size>) reg reg reg

**Pseudocode:**

Evaluate(WrEn);



```
for (n = 0; n < exec_size; n++) {  
  for (n = 0; n < exec_size; n++) {  
    float dwP = src0.RegNum.SubRegNum[bits4:2] // a dword aligned scalar  
    float dwQ = src0.RegNum.(SubRegNum[bit4:2]|0x1) // 2nd component  
    float dwR = src0.RegNum.(SubRegNum[bit4:2]|0x3) // 4th component  
    src2 = src1.(RegNum + 1) // Next GRF register  
    if (WrEn.chan[n] == 1) {  
      dst.chan[n] = dwP * src1.chan[n] + dwQ * src2.chan[n] + dwR  
    }  
  }  
}
```

#### Description:

The *pln* instruction computes a component-wise plane equation ( $w = p*u+q*v+r$  where  $u/v/w$  are vectors and  $p/q/r$  are scalars) of  $\langle src0 \rangle$  and  $\langle src1 \rangle$  and storing the results in  $\langle dst \rangle$ .  $\langle src1 \rangle$  is the input vector  $u$ . The second input vector  $v$  is implied from  $\langle src1 \rangle$ , as the next adjacent GRF register.  $\langle src0 \rangle$  provides input scalars  $p$ ,  $q$  and  $r$ , where  $p$  is the scalar value based on the region description of  $\langle src0 \rangle$  and  $q$  and  $r$  are the scalar values implied from  $\langle src0 \rangle$  region. Specifically,  $q/r$  is the second/fourth component of the 4-tuple (128-bit aligned) that  $p$  belongs to.

When *pln* instruction is used in SIMD16 form, the same input data channels  $p/q/r$  for  $\langle src0 \rangle$  are used for both SIMD8 instructions. However, as  $\langle src1 \rangle$  has two vectors  $u/v$ , where  $v$  is implied, the second SIMD8 instruction takes  $src1.(RegNum+2)$  as the second source operand.

#### Restrictions:

This is a specialized instruction that only support execution size of 8 or 16.

$\langle src0 \rangle$  region must be a replicated scalar (with  $HorzStride = VertStride = 0$ ).

$\langle src0 \rangle$ ,  $\langle src1 \rangle$ ,  $\langle src2 \rangle$ , and  $\langle dst \rangle$  may be float.

Source operands cannot be an accumulator register.

$\langle src0 \rangle$  for *pln* instruction has to has .0 or .4 as the subregister number.

[DevSNB]: $\langle src1 \rangle$  must be even register aligned.



### 8.3.36 ret – Return [DevSNB]

Opcode	Instruction	Description
45 (0x2d)	ret NULL <src0>	Return back from a subroutine.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•				DWord	

#### Format:

[(<pred>)] ret (<exec\_size>) NULL <src0>

#### Syntax:

[(<pred>)] ret (<exec\_size>) NULL reg

#### Pseudocode:

```
Evaluate(WrEn);
for (n = 0; n < 16; n++) {
    if (WrEn.channel[n] == 1) {
        PclIP[n] = src0.chan[0];
        CallMask[n] = 0;
    } else {
        PclIP[n] = IP + 1;
    }
}

if (CallMask[n:0] == 0) { // all channels are zero
    Jump(src0.chan[0]);
    CallMask = src0.chan[1];
}
```

#### Description:

The *ret* instruction returns the execution from a subroutine back.

The *ret* instruction can be predicated or non-predicated. If non-predicated, all the channels will jump to the return IP in the first channel of *src0*. The *CallMask* will be restored from the second channel of *src0*. If predicated, the enabled channels will jump to the return IP in the first channels of *src0*, the corresponding



bit in the CallMask will be reset to zero, if all the bits in the CallMask are zero after the *call* instruction, the execution will jump to the return IP in the first channel of src0.

If SPF is ON, none of the PclIP is updated.

When SPF is on, the predication control must be scalar.

**Restrictions:**

This instruction cannot take accumulator as source.

### 8.3.37 rddd – Round Down

Opcode	Instruction	Description
69 (0x45)	rddd <dst> <src0>	Taking component-wise floating point downward rounding of <src0> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]

**Format:**

[(<pred>)] rddd[.<cmod>] (<exec\_size>) <dst> <src0>

**Syntax:**

[(<pred>)] rddd[.<cmod>] (<exec\_size>) reg reg  
 [(<pred>)] rddd[.<cmod>] (<exec\_size>) reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = floor(src0.chan[n]);
    }
}
```

**Description:**

The *rddd* instruction takes component-wise floating point downward rounding (to the integral float number closer to negative infinity) of <src0> and storing the rounded integral float results in <dst>. This is commonly referred to as the *floor()* function.

This instruction only applies to floating point source and destination operands.



Output data <dst> for floating point rounding-down follow rules in Table 8-7 (or Table 8-8), if the current floating point mode is IEEE mode (or ALT mode).

**Table 8-7. Floating point round-down in IEEE mode**

<src0>	-inf	-finite	-denorm	-0	+0	+denorm	+finite	+inf	NaN
<dst>	-inf	*	-0	-0	+0	+0	**	+inf	NaN
Notes:									
* Result may be {-finite, -0}.									
** Result may be {+finite, +0}.									

**Table 8-8. Floating point round-down in ALT mode**

<src0>	- fmax	-finite	-denorm	-0	+0	+denorm	+finite	+ fmax	***
<dst>	-fmax	*	-0	-0	+0	+0	**	+fmax	
Notes:									
* Result may be {-finite, -0}.									
** Result may be {+finite, +0}.									
*** Result is undefined if <src0> is {-inf, +inf, NaN}.									

**Restrictions:**

This instruction cannot take accumulator as source or destination operand. However, when the accumulator is implicitly updated by this instruction, the results in the accumulator are undefined.



### 8.3.38 rndu – Round Up

Opcode	Instruction	Description
68 (0x44)	<code>rndu &lt;dst&gt; &lt;src0&gt;</code>	Taking component-wise floating point upward rounding of <src0> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]

**Format:**

`[(<pred>)] rndu[.<cmode>] (<exec_size>) <dst> <src0>`

**Syntax:**

`[(<pred>)] rndu[.<cmode>] (<exec_size>) reg reg`  
`[(<pred>)] rndu[.<cmode>] (<exec_size>) reg imm32`

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        if (src0.chan[n]-floor(src0.chan[n]) > 0.0f)
            dst.chan[n] = floor(src0.chan[n]) + 1;
        else
            dst.chan[n] = src0.chan[n];
    }
}
```

**Description:**

The *rndu* instruction takes component-wise floating point upward rounding (to the integral float number closer to positive infinity) of <src0>, commonly known as the *ceiling()* function.

This instruction only applies to floating point source and destination operands.





**Table 8-9. Floating point round-up in IEEE mode**

<src0>	-inf	-finite	-denorm	-0	+0	+denorm	+finite	+inf	NaN
<dst>	-inf	*	-0	-0	+0	+0	**	+inf	NaN
Notes: * Result may be {-finite, -0}. ** Result may be {+finite, +0}. *** Increment may be {0, 1}. It is 0 if source data is an integral float, and is 1 otherwise.									

**Table 8-10. Floating point round-up in ALT mode**

<src0>	- fmax	-finite	-denorm	-0	+0	+denorm	+finite	+ fmax	****
<dst>	-fmax	*	-0	-0	+0	+0	**	+fmax	
Notes: * Result may be {-finite, -0}. ** Result may be {+finite, +0}. *** Increment may be {0, 1}. It is 0 if source data is an integral float, and is 1 otherwise. **** Result is undefined if <src0> is {-inf, +inf, NaN}.									



### 8.3.39 rnde – Round to Even

Opcode	Instruction	Description
70 (0x46)	Rnde <dst> <src0>	Taking component-wise floating point round-to-even operations of <src0> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]

**Format:**

[(<pred>)] rnde[.<cmod>] (<exec\_size>) <dst> <src0>

**Syntax:**

[(<pred>)] rnde[.<cmod>] (<exec\_size>) reg reg  
 [(<pred>)] rnde[.<cmod>] (<exec\_size>) reg imm32

**Pseudocode:**

```

Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = floor(src0.chan[n]);
        if (src0.chan[n]-floor(src0.chan[n]) > 0.5f) {
            dst.chan[n] = floor(src0.chan[n]) + 1;
        } else if (src0.chan[n]-floor(src0.chan[n]) < 0.5f) {
            dst.chan[n] = floor(src0.chan[n]);
        } else {
            if (dst.chan[n] is odd) {
                dst.chan[n] = floor(src0.chan[n]) + 1;
            } else {
                dst.chan[n] = floor(src0.chan[n]);
            }
        }
    }
}
  
```

**Description:**

The *rnde* instruction takes component-wise floating point round-to-even operation of <src0> with results in two pieces – a downward rounded integral float results stored in <dst> and the round-to-even increments



stored in the rounding increment bits. The round-to-even increment must be added to the results in <dst> to create the final round-to-even values to emulate the round-to-even operation, commonly known as the *round()* function. The final results are the one of the two integral float values that is nearer to the input values. If the neither possibility is nearer, the even alternative is chosen.

Output data <dst> for floating point rounding-to-even follow rules in Table 8-11 (or Table 8-12), if the current floating point mode is IEEE mode (or ALT mode).

**Table 8-11. Floating point round-to-even in IEEE mode**

<src0>	-inf	-finite	-denorm	-0	+0	+denorm	+finite	+inf	NaN
<dst>	-inf	*	-0	-0	+0	+0	**	+inf	NaN
Notes:									
* Result may be {-finite, -0}.									
** Result may be {+finite, +0}.									
*** Increment may be {0, 1}. It is 0 if source data is an integral float. It may be 0 or 1 otherwise.									

**Table 8-12. Floating point round-to-even in ALT mode**

<src0>	- fmax	-finite	-denorm	-0	+0	+denorm	+finite	+ fmax	****
<dst>	-fmax	*	-0	-0	+0	+0	**	+fmax	
Notes:									
* Result may be {-finite, -0}.									
** Result may be {+finite, +0}.									
*** Increment may be {0, 1}. It is 0 if source data is an integral float. It may be 0 or 1 otherwise.									
**** Result is undefined if <src0> is {-inf, +inf, NaN}.									



### 8.3.40 *rndz* – Round to Zero

Opcode	Instruction	Description
71 (0x47)	<i>rndz</i> <dst> <src0>	Taking component-wise floating point round-to-zero operations of <src0> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT]	[FLT]

**Format:**

[[<pred>]] *rndz*.[<cmod>] (<exec\_size>) <dst> <src0>

**Syntax:**

[[<pred>]] *rndz*.[<cmod>] (<exec\_size>) reg reg  
 [[<pred>]] *rndz*.[<cmod>] (<exec\_size>) reg imm32

**Pseudocode:**

```

Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = floor(src0.chan[n]);
        if (abs(src0.chan[n]) < abs(dst.chan[n])) {
            dst.chan[n] = floor(src0.chan[n]) + 1;
        } else {
            dst.chan[n] = floor(src0.chan[n]);
        }
    }
}

```

**Description:**

The *rndz* instruction takes component-wise floating point round-to-zero operation of <src0> with results in two pieces – a downward rounded integral float results stored in <dst> and the round-to-zero increments stored in the rounding increment bits. The round-to-zero increment must be added to the results in <dst> to create the final round-to-zero values to emulate the round-to-zero operation, commonly known as the *truncate()* function. The final results are the one of the two closest integral float values to the input values that is nearer to zero.

Output data <dst> for floating point rounding-to-zero follow rules in Table 8-13 (or Table 8-14), if the current floating point mode is IEEE mode (or ALT mode).



**Table 8-13. Floating point round-to-zero in IEEE mode**

<src0>	-inf	-finite	-denorm	-0	+0	+denorm	+finite	+inf	NaN
<dst>	-inf	*	-0	-0	+0	+0	**	+inf	NaN
Notes: * Result may be {-finite, -0}. ** Result may be {+finite, +0}. *** Increment may be {0, 1}. It is 0 if source data is an integral float, and is 1 otherwise.									

**Table 8-14. Floating point round-to-zero in ALT mode**

<src0>	- fmax	-finite	-denorm	-0	+0	+denorm	+finite	+ fmax	****
<dst>	-fmax	*	-0	-0	+0	+0	**	+fmax	
Notes: * Result may be {-finite, -0}. ** Result may be {+finite, +0}. *** Increment may be {0, 1}. It is 0 if source data is an integral float, and is 1 otherwise. **** Result is undefined if <src0> is {-inf, +inf, NaN}.									



### 8.3.41 sad2 – Sum of Absolute Difference 2

Opcode	Instruction	Description
80 (0x50)	sad2 <dst> <src0> <src1>	Performing a two-wide sum-of-absolute-difference operation on a 2-tuple basis of <src0> and <src1>, and storing the scalar result to the first channel per 2-tuple in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[INT]	[INT]

**Format:**

[(<pred>)] sad2[.<cmode>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] sad2[.<cmode>] (<exec\_size>) reg reg reg  
 [(<pred>)] sad2[.<cmode>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```

Evaluate(WrEn);
for (n = 0; n < exec_size; n+=2) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = abs(src0.chan[n] - src1.chan[n])
            + abs(src0.chan[n+1] - src1.chan[n+1]);
    }
}
  
```

**Description:**

The *sad2* instruction takes source data channels from <src0> and <src1> in groups of 2-tuples. For each 2-tuple, it computes the sum-of-absolute-difference (SAD) between <src0> and <src1> and stores the scalar result in the first channel of the 2-tuple in <dst>.

This instruction only applies to integer operands. In particular, source operands must be unsigned bytes and/or signed bytes and destination operand must be of word type. Source modifiers are allowed.

The results are also stored in the accumulator register. Destination operand and accumulator maintain 16-bit per channel precision.

Destination register must have a stride of 2 words and must be aligned to even word. The even words in destination region will contain the correct data. The odd words are also written but with undefined values

**Restrictions:**

Source operands cannot be an accumulator register.

Execution size cannot be 1 as the computation requires at least two data channels.



### 8.3.42 sada2 – Sum of Absolute Difference Accumulate 2

Opcode	Instruction	Description
81 (0x51)	sada2 <dst> <src0> <src1>	Performing a two-wide sum-of-absolute-difference operation on a 2-tuple basis of <src0> and <src1>, added to that from the accumulator, and storing the scalar result to the first channel per 2-tuple in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[INT]	[INT]

**Format:**

[(<pred>)] sada2[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] sada2[.<cmod>] (<exec\_size>) reg reg reg  
 [(<pred>)] sada2[.<cmod>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n+=2) {
    uwTmp = abs(src0.channel[n] - src1.channel[n])
           + abs(src0.channel[n+1] - src1.channel[n+1])
    if (WrEn.channel[n] == 1) {
        dst.channel[n] = uwTmp + acc[n]
    }
}
```

**Description:**

The *sada2* instruction takes source data channels from <src0> and <src1> in groups of 2-tuples. For each 2-tuple, it computes the sum-of-absolute-difference (SAD) between <src0> and <src1>, adds the intermediate result with the accumulator value corresponding to the first channel, and stores the scalar result in the first channel of the 2-tuple in <dst>.

This instruction only applies to integer operands. In particular, source operands must be unsigned bytes and/or signed bytes and destination operand must be of word type. Source modifiers are allowed.

The results are also stored in the accumulator register. Destination operand and accumulator maintain 16-bit per channel precision. Higher precision (guard bits) stored in the accumulator allows multiple rounds (64 rounds) of *sada2* instructions to be issued back to back without overflow the accumulator.



Destination register must have a stride of 2 words and must be aligned to even word. The even words in destination region will contain the correct data. The odd words are also written but with undefined values

**Restrictions:**

Source operands cannot be an accumulator register.

Execution size cannot be 1 as the computation requires at least two data channels.

### 8.3.43 sel – Select

Opcode	Instruction	Description
2 (0x02)	(pred) sel <dst> <src0> <src1>	Component-wise selective move from <src0> or <src1> to <dst> based on predication or cmod result. The sel instruction can not use accumulator as destination

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[FLT] [INT]	[FLT] [INT]

**Format:**

(<pred>) sel[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

(<pred>) sel[.<cmod>] (<exec\_size>) reg reg reg  
 (<pred>) sel[.<cmod>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```

Evaluate(WrEn, NoPMask);

if (cmod == "0000") { // no CMod
  Evaluate(PMask);
  for (n = 0; n < exec_size; n++) {
    if (WrEn.channel[n] == 1) {
      if (PMask.channel[n] == 1) {
        dst.channel[n] = src0.channel[n]
      } else {
        dst.channel[n] = src1.channel[n]
      }
    }
  }
}
  
```





```
    }
    else { // with CMod
        Evaluate(CMod);
        for (n = 0; n < exec_size; n++) {
            if (WrEn.channel[n] == 1) {
                if (CMod.channel[n] == 1) {
                    dst.channel[n] = src0.channel[n]
                } else {
                    dst.channel[n] = src1.channel[n]
                }
            }
        }
    }
}
```

### Description:

The *sel* instruction selectively moves the components in <src0> or <src1> into the channels of <dst> based on the predication. On a channel by channel basis, if the channel condition is true, data in <src0> is moved into <dst>; Otherwise, data in <src1> is moved into <dst>.

As the predication is used to select the two sources, it is not included in the evaluation of WrEn. <pred> is mandatory if <cmode> is "0000". If it is <omitted> and <cmode> is "0000", the results are unpredictable.

In case of <cmode> not equal to "0000", a compare will be performed and the result flag will be used for the *sel* instruction. <cmode> .ge and .l follow the CMPN rules, and all other <cmode> follow the CMP rules. <pred> is not allowed in this mode.

*sel* instruction with <cmode> .l should be used to emulate MIN instruction.

*sel* instruction with <cmode> .ge should be used to emulate MAX instruction.

If any of the source of the *sel* instruction is NaN, the non-NaN source will be the result, in case both sources are NaN, the result will also be NaN. This only applies to *sel* instruction with .l and .ge conditional modifier. For the other conditional modifiers, src1 will be always be selected if either or both sources is/are NaN.

*sel* instruction with <cmode> will flush denorm to zero; *sel* instruction without <cmode> will retain denorm.

Format conversion is not allowed.

[DevSNB]: *sel* instruction can not use accumulator source.

### Restrictions:

Destination channels cannot be on odd-byte sub-register locations. In other words, when destination is of byte type, destination horizontal stride cannot be 1. If destination horizontal stride is not 1, destination register region origin cannot be on an odd byte location. This is because that the conditional flag for execution channels that have minimal granularity of word are used by this instruction.

[DevSNB]: The *sel* instruction can not use accumulator as destination



### 8.3.44 send – Send Message [DevSNB]

Opcode	Instruction	Description
49 (0x31)	send <dest> <src> <ex_desc> <desc>	Send a message stored in MRF starting at <src> to a shared function identified by <ex_desc> along with control from <desc> with a GRF writeback location at <dest>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•					[FLT] [INT]

**Format:**

[[<pred>]] send (<exec\_size>) <dest> <src> <ex\_desc> <desc>

**Syntax:**

[[<pred>]] send (<exec\_size>) reg mreg imm6 reg32a

[[<pred>]] send (<exec\_size>) reg mreg imm6 imm32

**Pseudocode:**

```
Evaluate(WrEn);
<MsgChEnable> = WrEn;
<SourceReg> = <src>.RegNum;
MessageQueue(<MsgChEnable>, <ResponseReg>, <SourceReg>, <desc>,
<ex_dest>);
```

**Description:**

The *send* instruction performs data communication between a thread and external function units, including shared functions (Sampler, Data Port Read, Data Port Write, URB, and Message Gateway) and some fixed functions (e.g. Thread Spawner, who also have an unique Shared Function ID). The *send* instruction adds an entry to the EU's message request queue. The request message is stored in a block of contiguous MRF registers. The response message, if present, will be returned to a block of contiguous GRF registers. The return GRF writes may be in any order depending on the external function units. <src> is the lead MRF register for request. <dest> is the lead GRF register for response. The message descriptor field <desc> contains the Message Length (the number of consecutive MRF registers) and the Response Length (the number of consecutive GRF registers). It also contains the header present bit, and the function control signals. The extend message descriptor field <ex\_desc> contains the target function ID. WrEn is forwarded to the target function in the message sideband.

The *send* instruction is the only way to terminate a thread. When the EOT (End of Thread) bit of <desc> is set, it indicates the end of thread to the EU, the Thread Dispatcher and, in most cases, the parent fixed function.



Message descriptor field <desc> can be a 32-bit immediate, *imm32*, or a 32-bit scalar register, <reg32a>. DevSNB restricts that the 32-bit scalar register <reg32a> must be the leading dword of the address register. It should be in the form of a0.0<0;1,0>:ud. When <desc> is a register operand, only the lower 29 bits of <reg32a> are used.

<ex\_desc> is a 6-bit immediate, *imm6*. The lower 4bits of the <ex\_desc> specifies the SFID for the message. The MSb of the message descriptor, the EOT field, always comes from bit 127 of the instruction word, which is the MSb of *imm6*. A thread must terminate with a *send* instruction with EOT turned on.

<src> is a 256-bit aligned MRF register. It serves as the leading MRF register of the request.

<dest> serves for two purposes: to provide the leading GRF register location for the response message if present, and to provide parameters to form the channel enable sideband signals.

<dest> signals whether there is a response to the message request. It can be either a null register, a direct-addressed GRF register or a register-indirect GRF register. Otherwise, hardware behavior is undefined.

If <dest> is null, there is no response to the request. Meanwhile, the Response Length field in <desc> must be 0. Certain types of message requests, such as memory write (store) through the Data Port, do not want response data from the function unit. If so, the posted destination operand can be null.

If <dest> is a GRF register, the register number is forwarded to the shared function. In this case, the target function unit must send one or more response message phases back to the requesting thread. The number of response message phases must match the Response Length field in <desc>, which of course cannot be zero. For some cases, it could be an empty return message. An empty return message is defined as a single phase message with all channel enables turned off.

The subregister number, horizontal stride, destination mask and type fields of <dest> are always valid and are used in part to generate the WrEn. This is true even if <dest> is a null register (this is an exception for null as for most cases these fields are ignored by hardware). These parameters of <dest> follow the same restriction as that of normal destination operand – destination region cannot cross the 256-bit register boundary.

The 16-bit channel enables of the message sideband are formed based the WrEn. Interpretation of the channel enable sideband signals is subject to the target external function. In general for a 'send' instruction with return messages, they are used as the destination dword write mask for the GRF registers starting at <dest>. For a message that has multiple return phases, the same set of channel enable signals applies to all the return phases.

Thread managed memory coherency: A special usage of using non-null <dest> is to support write-commit signaling for memory write service by the Data Port Write unit. If <post\_dest> is not null for a memory write request, the Data Port along with the Data Cache or Render Cache will wait until all the posted writes for the request have reached the coherent domain before sending back to the requesting thread an empty message to <dest> register. A memory write reaching the coherent domain, also referred to as reaching the global observable state, means that subsequent read to the same memory location, no matter which thread issues the read, must return the data of the write.

The destination dependency control, {NoDDClr}, can be used in this instruction. This allows software to control the destination dependencies for multiple 'read'-type messages similar to that for multiple instructions using EU execution pipeline. As *send* does not check register dependencies for the post destination, {NoDDChk} should not be used for this instruction.



DevSNB supports both header and headerless messages, and it's indicated by the "Header Present" bit in the message descriptor.

**Table 8-15. Message Descriptor Definition**

Bit	Description
31:29	Reserved : MBZ
28:25	<p><b>Message Length.</b> This field specifies the number of 256-bit MRF registers starting from &lt;curr_dest&gt; to be sent out on the request message payload. Valid value ranges from 1 to 15. A value of 0 is considered erroneous.</p> <p>Format = U4</p> <p>Range = [1,15]</p>
24:20	<p><b>Response Length.</b> This field indicates the number of 256-bit registers expected in the message response. The valid value ranges from 0 to 16. A value 0 indicates that the request message does not expect any response. The largest response supported is 15 GRF registers.</p> <p>Format = U5</p> <p>Range = [0,16]</p>
19	<p><b>Header Present.</b> If set, indicates that the message includes a header. Depending on the target shared function, this field may be restricted to either enabled or disabled. Refer to the specific shared function section for details.</p> <p>Format = Enable</p>
18:0	<p><b>Function Control</b></p> <p>This field is intended to control the target function unit. Refer to the section on the specific target function unit for details on the contents of this field.</p>

**Table 8-16. Extended Message Descriptor Definition**

Bit	Description
5	<p><b>End Of Thread</b></p> <p>This field, if set, indicates that this is the final message of the thread and the thread's resources can be reclaimed.</p>
4	reserved
3:0	<p><b>Target Function ID</b></p> <p>This field indicates the function unit for which the message is intended.</p> <p>0000: Null</p> <p>0001: Reserved</p>



Bit	Description
	0010: Sampling Engine 0011: Message Gateway 0100: Data Port Sampler Cache 0101: Data Port Render Cache 0110: URB 0111: Thread Spawner 1000: Video Motion Estimation 1001: Constant Cache 1010 – 1111: Reserved

The following table provides a summary of the signals associated with each message that is sent to the shared function. It contains fields from the send instruction as well as fields from control register *cr0* and state register *sr0*.



**Table 8-17. Sideband Signals Associated with Each Message Sent to the Shared Function**

Signal	Bits	Source
EOT	1	End of Thread: Sourced from the EOT bit in <i>send</i> instruction word
SFID	3	Shared Function Identifier: Sourced from the target function ID field in <ex_desc> of <i>send</i>
MLEN	4	Message Length: Sourced from the message length field in <desc> of <i>send</i>
RLEN	5	Response Length: Sourced from the response length field in <desc> of <i>send</i>
FC	19	Function Control: Sourced from the function control field in <desc> of <i>send</i>
REG	7	Destination Register: Sourced from the 256-bit register aligned register number of the <dest> field of <i>send</i>
CE	16	Channel Enable: Sourced from the write enable of <i>send</i>
CLEAR	1	Destination Register Clear: Source from the Destination Dependency Control field (inverse of NoDDClr) in <i>send</i> instruction word
FFID	4	Fixed Function Identifier: Sourced from the Fixed Function ID field in <i>sr0</i>
EUID	4	Execution Unit Identifier: Sourced from the EUID field in <i>sr0</i>
TID	2	Thread Identifier: Sourced from the TID field in <i>sr0</i>

**Restrictions:**

Software must obey the following rules in signaling the end of thread using the *send* instruction:

- The posted destination operand must be null.
  - No acknowledgement is allowed for the *send* instruction that signifies the end of thread. This is to avoid deadlock as the EU is expecting to free up the terminated thread's resource.
- A thread must terminate with a *send* instruction with message to a shared function on the output message bus; therefore, it cannot terminate with a *send* instruction with message to the following shared functions: Sampler unit, NULL function
- For example, a thread may terminate with a URB write message or a render cache write message.
- A root thread originated from the media (generic) pipeline must terminate with a *send* instruction with message to the Thread Spawner unit. A child thread should also terminate with a *send* to TS. Please refer to the Media Chapter for more detailed description.

The *send* instruction can not update accumulator registers.

Saturate is not supported for *send* instruction.

ThreadCtrl are not supported for *send* instruction.



R0 or NULL register can not be used as destination for all messages with writebacks

**[Errata: DevSNB(SNB)]** If MRF register is updated by any instruction that “indexed/indirect” source AND is followed by a *send*, the instruction requires a “Switch”. This is to avoid race condition where *send* may dispatch before MRF is updated.

### 8.3.45 sendc – Conditional Send Message

Opcode	Instruction	Description
49 (0x31)	sendc <dest> <src> <desc> <ex_dest>	Wait for the dependencies in the TDR register cleared, then send a message stored in MRF starting at <str> to a shared function identified by <ex_desc> along with control from <desc> with a GRF writeback location at <dest>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•					[FLT] [INT]

**Format:**

[(<pred>)] sendc (<exec\_size>) <dest> <src> <exdesc> <desc>

**Syntax:**

[(<pred>)] sendc (<exec\_size>) reg reg reg32a imm4  
[(<pred>)] sendc (<exec\_size>) reg reg imm32 imm4

**Pseudocode:**

```

if ((TDR[7]... || TDR[2] || TDR[1] || TDR[0]) == TRUE) {
    wait;
}
Evaluate(WrEn);
<MsgChEnable> = WrEn;
<SourceReg> = <src>.RegNum;
MessageQueue(<MsgChEnable>, <ResponseReg>, <SourceReg>, <desc>,
<ex_dest>);

```

**Description:**

The *sendc* instruction has the same behavior as the *send* instruction except the following.

The *sendc* instruction first check the dependent threads inside the Thread Dependency Register, there are up to 4 dependent threads in the TDR register. The *sendc* instruction will be executed only when all the dependent threads in the TDR register are retired.



**Restrictions:**

The *sendc* instruction has the same restrictions as the *send* instruction.

### 8.3.46 shl – Shift Left

Opcode	Instruction	Description
9 (0x09)	shl <dst> <src0> <src1>	Performing component-wise logic left shift of <src0> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[INT]	[INT]

**Format:**

[(<pred>)] shl[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[(<pred>)] shl[.<cmod>] (<exec\_size>) reg reg reg  
 [(<pred>)] shl[.<cmod>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```

Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = src0.chan[n] << src1.chan[n]
    }
}
  
```

**Description:**

The *shl* instruction performs component-wise logical left shift of <src0> with zero-insertion and storing the results in <dst>. The amount of bit shift is provided by <src1>, where only the 5 LSBs of each channel of <src1> are used as an unsigned integer value. The MSBs of <src1> data channels are ignored. . The destination operand cannot be the accumulator.

5-bit shifting applies to packed-dword mode and packed-word mode. For packed word mode, the accumulators have 34 bits per channel. <src0> and <dst> can be signed or unsigned integers and can be of different types. This instruction does not work with float type operands. Saturation modifier is only allowed when this instruction is in packed-word mode. Hardware detects overflow properly and use it to perform saturation operation on the output, as long as the shifted result is within 34 bits. Otherwise, the result is undefined.

Results of saturation in packed-dword mode are unpredictable.



**Restrictions:**

This instruction does not work with float type operands.

### 8.3.47 shr – Shift Right

Opcode	Instruction	Description
8 (0x08)	shr <dst> <src0> <src1>	Performing component-wise logic right shift of <src0> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•	•	•	•	[INT]	[INT]

**Format:**

[[<pred>]] shr[.<cmod>] (<exec\_size>) <dst> <src0> <src1>

**Syntax:**

[[<pred>]] shr[.<cmod>] (<exec\_size>) reg reg reg  
[[<pred>]] shr[.<cmod>] (<exec\_size>) reg reg imm32

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = src0.chan[n] >> src1.chan[n]
    }
}
```

**Description:**

The *shr* instruction performs component-wise logical right shift of <src0> with zero-insertion and storing the results in <dst>. The amount of bit shift is provided by <src1> where only the 5 LSBs of each channel of <src1> are used as an unsigned integer value. The MSBs of <src1> data channels are ignored.

5-bit shifting applies to packed-dword mode and packed-word mode. For packed word mode, the accumulators have 33 bits per channel.

This instruction only takes on unsigned sources. When <src0> contains unsigned integers, no source modifier is allowed. <src0> is only allowed to be signed integer if source modifier (*abs*) is used. *Note: for unsigned sources, the behavior of shr and asr are effectively the same.*

**Restrictions:**

This instruction does not work with float type operands.



### 8.3.48 wait – Wait Notification

Opcode	Instruction	Description
48 (0x30)	wait <nreg>	Waiting for notification on the notification register <nreg>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types

**Format:**

wait (<exec\_size>) <nreg>

**Syntax:**

wait (1) n#

**Pseudocode:**

n/a

**Description:**

The *wait* instruction evaluates the value of the notification count register <nreg>. If <nreg> is zero, the execution of the thread is stalled and the thread is put in 'wait\_for\_notification' state. If <nreg> is not zero (i.e., one or more notifications have been received), <nreg> is decremented by one and the thread continues executing on the next instruction. If a thread is in the 'wait\_for\_notification' state, when a notification arrives, the notification count register is incremented by one. As the notification count register becomes non-zero, the thread wakes up to continue execution and at the same time the notification register is decremented by one. If there was only one notification arrived, the notification register value becomes zero. However, during the above mentioned time period, it is possible that more notifications may arrive, making the notification register non-zero again.

When multiple notifications are received, software must use 'wait' instruction to decrement notification count register for each of the notifications.

Notification register n0:ud is for thread to thread communication (through message gateway shared function) and n1:ud for host to thread communication (through MMIO registers). See Message Gateway chapter for thread-thread communication.

**Restrictions:**

Only one source operand.

<src0> and <dst> must be n0, n1, or n2, <src1> must be *null*.

Execution size must be 1 as the notification registers are scalar.

Predication is not allowed.



### Implementation restriction:

Two back-to-back *wait* instructions in a program (without any instruction in between) are not allowed. As a minimal, a *nop* has to be inserted between two *wait* instructions.

## 8.3.49 while – While [DevSNB]

Opcode	Instruction	Description
39 (0x27)	while <dst> <src0> <src1> <JIP>	Marking the end of a do-while block of code.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•					

### Format:

```

[(<pred>)] while (<exec_size>) null null null <JIP>
while (<exec_size>) null <src0> <src1> <JIP>

```

### Syntax:

```

[(<pred>)] while (<exec_size>) null null null imm16
while (<exec_size>) null reg reg imm16
while (<exec_size>) null reg imm32 imm16

```

### Pseudocode:

```

Evaluate(WrEn);
for (n = 0; n < 16; n++) {
    if (WrEn.channel[n] == 1) {
        if (<cmode> == 0) { // no embedded compare
            if (PMask.channel[n] == 1) {
                PcIP[n] = IP + <JIP>;
            }
        }
        else {
            PcIP[n] = IP + 1;
        }
    }
    else { // with embedded compare
        if (cmode.channel[n] == 1) {
            PcIP[n] = IP + <JIP>;
        }
        else {

```



```

                                PciIP[n] = IP + 1;
                                }
                            }
                    }
}
if (<cmode> == 0) { // no embedded compare
    if (!PMask == 1) { // any enabled channel true
        Jump(IP + <JIP>);
    }
}
else { // with embedded compare
    if (!cmode == 1) { // any enabled channel true
        Jump(IP + <JIP>);
    }
}
}

```

**Description:**

The *while* instruction marks the end of a do-while block. The instruction first evaluates the loop termination condition for each channel based the current channel enables the predication flag specified in the instruction. If any channel has not terminated, a branch is taken to a destination address based specified in the instruction, and the loop continued for those channels. Otherwise, execution continue down to the next instruction.

The following table describes the 16-bit jump target offset <JIP>. <JIP> is a signed 16-bit number, added to IP pre-increment, and should point to the first instruction with the *do* label of the do-while block of code. It should be a negative number for the backward referencing. In DevSNB binary, <JIP> is at location <dst> and must be of type W (signed word integer).

Bit	Description
15:0	<b>JIP (Jump Target Offset).</b> This field specifies the jump distance in number of 64bits data chunks if a jump is taken for the instruction. Format = S15. Signed integer in 2's compliment

If SPF is ON, none of the PciIP is updated.

**Restrictions:**

To use embedded compares, the predicate control field for this instruction must be zero, and the conditional modifier field must be none zero. The destination must follow the rules below:

5. Must has the same element size as source0
6. Must have horizontal stride of 1

To use predicated *if* instruction, the conditional modifier field must be zero.



### 8.3.50 xor – Logic Xor

Opcode	Instruction	Description
7 (0x07)	xor <dst> <src0> <src1>	Performing component-wise logic XOR of <src0> and <src1> and storing the results in <dst>.

Pred	Sat	Cond Mod	Src Mod	Src Types	Dst Types
•		•	•	[INT]	[INT]

**Format:**

`[(<pred>)] xor[.<cmod>] (<exec_size>) <dst> <src0> <src1>`

**Syntax:**

`[(<pred>)] xor[.<cmod>] (<exec_size>) reg reg reg`  
`[(<pred>)] xor[.<cmod>] (<exec_size>) reg reg imm32`

**Pseudocode:**

```
Evaluate(WrEn);
for (n = 0; n < exec_size; n++) {
    if (WrEn.chan[n] == 1) {
        dst.chan[n] = src0.chan[n] ^ src1.chan[n];
    }
}
```

**Description:**

The *xor* instruction performs component-wise logic XOR operation between <src0> and <src1> and stores the results in <dst>.

Source modifiers are allowed.

Accumulator register is allowed to be the destination of this instruction with the restrictions listed below.

**Restrictions:**

Sign (SN) and Overflow (OF) conditions are undefined for this logic instruction. Consequently, saturation modifier (.sat) is not allowed.

This instruction does not work with float type operands.

Accumulator is allowed to be the source of this instruction, but source modifier is not allowed for an accumulator source.



When accumulator is the destination of this instruction, only the low bits corresponding to the data type (16 bits for word or 32 bits for dword integer instruction) in the accumulator contain the correct results. The internal extra-precision bits as well as the sign bit of the accumulator are undefined. Consequently, there are restrictions for subsequent instructions that use the data in the accumulator register created from the previous logical instruction.

- Only logical and data move instructions are allowed to source the accumulator. Results of other instructions (e.g. arithmetic or shift) are undefined.
- When the accumulator is the source of a data move (mov or sel) instruction, the destination operand must be of integer type (e.g. no conversion to float) and this instruction cannot have saturation instruction modifier.



## 9. EU Programming Guide

### 9.1 Assembler Pragmas

#### 9.1.1 Declarations

A register or a register region can be declared as a symbol using the following form

```
.declare <symbol>      Base=RegFile RegBase {.SubRegBase} ElementSize=ElementSize  
{SrcRegion=DefaultSrcRegion} {DstRegion=DefaultDstRegion} {Type=DefaultType}
```

The register file, the base of the register origin and the element size (in unit of bytes) are the mandatory parameters for a declared register region. Optionally, the base of the sub-register address, the default source region, the default destination region and the default type can be provided in the declaration for the symbol.

For immediate register addressing mode, the declared symbol can be used in the following Cartesian form

```
<symbol>(RegOff, SubRegOff)    ←    RegNum = RegBase + RegOff; SubRegNum =  
SubRegBase + SubRegOff
```

or in the following simplified row-aligned form

```
<symbol>(RegOff)    ←    RegNum = RegBase + RegOff; SubRegNum = SubRegBase
```

For register-indirect-register-addressing mode, the declared symbol can be used to provide immediate address term in the following Cartesian form

```
<symbol>[IdxReg, RegOff, SubRegOff]    ←    RegNum (byte-aligned) = [IdxReg] +(RegBase +  
RegOff)*32 + (SubRegBase + SubRegOff)*ElementSize
```

or in the following simplified row-aligned form

```
<symbol>[IdxReg, RegOff]    ←    RegNum (byte-aligned) = [IdxReg] +(RegBase + RegOff)*32
```

or in the form without the immediate address term

```
<symbol>[IdxReg]    ←    RegNum (byte-aligned) = [IdxReg] + RegBase
```



## 9.1.2 Defaults and Defines

The default execution size is set according to the destination register type as the following

Destination Register Type	Default Execution Size
UB   B	(16)
UW   W	(16)
F   UD   D	(8)

The default execution size can be overwritten globally for all instructions using

```
.default_execution_size (Execution_Size)
```

or be set according the **destination** register type using

```
.default_execution_size_Type (Execution_Size)
```

The default register type can be set for all register files using

```
.default_register_type Type
```

or be set per register file using

```
.default_register_type_RegFile Type
```

The default **source** register region for all symbols can be set using

```
.default_source_register_region <VirtStride; Width, HorzStride>
```

or be set per register type using

```
.default_source_register_region_type <VirtStride; Width, HorzStride>
```

The default **destination** register region for all symbols can be set using

```
.default_destination_register_region <HorzStride>
```

or be set per register type using

```
.default_destination_register_region_type <HorzStride>
```

Finally, the precompiler supports the string replacement statement of **.define** in the following form

```
.define <symbol> Expression
```







**Example 2: Declaration for 8x1 Float Regions:**

The following symbol *Trans* can be used to address any 8x1 float region within the Cartesian system of a 8x4 float GRF register area starting from r5.

Declaration	<pre>// 8x4 float Array starting at r5 .declare Trans Base=r5 ElementSize=4 Region=&lt;0;8,1&gt; Type=f</pre>
-------------	---

Fully-Expressed Instr	<pre>mov(8)      ?:f      r6.0&lt;0;8,1&gt;:f      // 2<sup>nd</sup> 16x1 Row of Trans. Matrix        // r5 FFFFFFFF        // r6 00000000        // r7 FFFFFFFF        // r8 FFFFFFFF</pre>
-----------------------	--

Short-handed Instr	<pre>mov      ?:f      Trans(1)      // RegNum = 5+1 = 6</pre>
--------------------	--

**Example 3: Declaration for 8x1 Float Regions with 1x1 Indirect Addressing:**

*Trans* region defined (same as in the previous example) is used in conjunction with the address register.

Declaration	<pre>//8x4 float data array and 16x1 word address array .declare Trans Base=r5 ElementSize=4 Region=&lt;0;8,1&gt; Type=f</pre>
-------------	--

Fully-Expressed Instr	<pre>mov(8)      ?:f      r[a0.0,224]&lt;0;8,1&gt;:f</pre>
-----------------------	--

Short-handed Instr	<pre>mov      ?:f      Trans[a0.0,2]      // [a0.0 + 5*32 + 2*32]</pre>
--------------------	---



**Example 4: Declaration with VxH Indirect Addressing:**

The VxH register-indirect-register-addressing for Trans can be provided in the following short-hand form.

Declaration	<pre>//8x4 float data array and word indices .declare Trans Base=r5 ElementSize=4 Region=&lt;0;8,1&gt; Type=f</pre>
Fully-Expressed Instr	<pre>mov(8)      ?:f      r[a0.0,224]&lt;1,0&gt;:f</pre>
Short-handed Instr	<pre>mov          ?:f          Trans[a0.0,2]&lt;1,0&gt;    // [a0.0+224] [a0.1+224] ... [a0.7+224]</pre>

**Example 5: Declaration with Vx1 Indirect Addressing:**

As width (4) is smaller than the execution region size (8), multiple indexed registers are used.

Declaration	<pre>//8x4 float data array and word address array .declare Trans Base=r5 ElementSize=4 Region=&lt;0;8,1&gt; Type=f</pre>
Fully-Expressed Instr	<pre>mov(8)      ?:f      r[a0.0,244]&lt;4,1&gt;:f</pre>
Short-handed Instr	<pre>mov          ?:f          Trans[a0.0,2]&lt;4,1&gt;    // [a0.0+224] [a0.1+224]</pre>



## 9.1.4 Assembly Programming Guideline

The following program skeleton illustrates the basic structure of a typical assembly program.

```
//      single line comment

/*
    block comment
*/

<preproc_directive>    // macros, include, etc. Are global – handled by the pre-processor
<preproc_directive>    // applies to all code that follows in sequence

// ----- some kernel -----
.kernel <kernel_name_string>    // [REQUIRED]

                                // ----- Register requirements -----
.reg_count_total    <uint>    // [REQUIRED] a more direct way to specify the exact parameters
require
.reg_count_payload    <uint>    // [REQUIRED] rather than to have to indirectly do that by adding the
                                // the payload and temps together to
get the total (as is the case now)

                                // Note: no more “reg-count-temp”

                                // ----- Defaults -----
<default...>    // these should be specified per-kernel and have
only kernel-scope
<default...>    // Same defaults as those already defined in the
ISA doc, but just
<default...>    // moved within the kernel to make each kernel
completely self-sufficient

                                // and not impacted defaults of earlier
kernels
```



```
placeholder for now...)
<MBDa>
<MBDb>
<MBDc>
<MBDd>

// ----- Memory Requirements -----
// [optional] memory block info (just a
// memory block descriptor a (TBD)
// memory block descriptor b (TBD)
// memory block descriptor c (TBD)
// memory block descriptor d (TBD)

// ----- Code -----
.code // [REQUIRED]
    <instruction>
    <instruction>
    <instruction>
    <LabelLine> // labels are code-block scope
    <instruction>
    <instruction>

.end_code // [REQUIRED]

.end_kernel // [REQUIRED]

// ----- next kernel -----

// ----- next kernel -----


// ...
```



## 9.2 Usage Examples

### 9.2.1 Vector Immediate

The immediate form of vector allows a constant vector to be in-lined in the instruction stream. An immediate vector is denoted by type *v* as *imm32:v*, where the 32-bit immediate field is partitioned into 8 4-bit subfields. Each 4-bit subfield contains a signed integer value in 2's complement form. Therefore each 4-bit subfield has a range of [-8, +7]. This is depicted in the following figure.



#### 9.2.1.1 Supporting OpenGL Vertex Shader Instruction SWZ

When an OpenGL Vertex Shader program is converted to run in Vertex Pair, i.e. two 4-wide vectors in parallel, the special OpenGL Shader instruction SWZ (Swizzle) needs to be emulated. OpenGL SWZ instruction uses an extended swizzle control field that, in addition to the 4-wide full swizzle control, also includes constant 0 and 1 replacement as well as per channel sign reversal. The later two are not supported by the native instruction. The vector immediate can significantly reduce the overhead of emulating such OpenGL instruction.

Consider an OpenGL Shader instruction in the form of

```
SWZ r1 r0.0-zx-1 // Expected results: r1.x = 0; r1.y = -r0.z; r1.z = r0.x; r1.w = -1
```

It can be emulated by the following three instructions.

```
mul (8) r1.0<1>:f r0.xzxz 0x1F111F11:v // Constant vector of (1 -1 1 1 1 -1 1 1)
```

```
mov (1) f0.0 8b'10011001 // Set flag & masked out channels y and z
```

```
(f0.0)mov(8) r1.0<1>:f 0x000F000F:v // Constant vector of (0 0 0 -1 0 0 0 -1)
```

In case that only 0, 1, -1 channel replacement is used and there is no signed swizzle, it may be emulated in two instructions. This is illustrated by the following example:

#### OpenGL:

```
SWZ r1 r0.0zx-1 // Expected results: r1.x = 0; r1.y = r0.z; r1.z = r0.x; r1.w = -1
```

#### DevSNB:

```
mov (1) f0.0 8b'01100110 // Set flag and masked out channels x and w
```

```
(f0.0)sel (8) r1.0<1>:f r0.yzxy 0x000F000F:v // Constant vector of (0 0 0 -1 0 0 0 -1)
```



## 9.2.2 Destination Mask for DP4 and Destination Dependency Control

The following example demonstrates the use of destination mask mode of floating point dot-product instruction as well as the use of destination dependency control to improve performance (i.e., avoiding unnecessary thread switch due to possible false dependencies).

With destination dependency control, the DevSNB instructions are as the following. The first instruction in the sequence checks for the destination dependency, but does not clear the dependency bit. The subsequent two instructions would do neither of them. The last instruction avoids checking the destination dependency, but at completion, it clears the destination scoreboard. It ensures that the content of the destination register is coherent, if any of the following instructions uses the same register as source.

```
dp4 (8) r5.0<1>.x:f r0.0<4;4,1>:f r4.0<4;4,1>:f {NoDDClr}
dp4 (8) r5.0<1>.y:f r1.0<4;4,1>:f r4.0<4;4,1>:f {NoDDClr, NoDDCChk}
dp4 (8) r5.0<1>.z:f r2.0<4;4,1>:f r4.0<4;4,1>:f {NoDDClr, NoDDCChk}
dp4 (8) r5.0<1>.w:f r3.0<4;4,1>:f r4.0<4;4,1>:f {NoDDChk}
```

Just as a comparison, IF DP4 implies reduction at the destination; additional shifted moves are required to achieve the same results. The corresponding codes are as the following. The lower performance due to the additional three move instruction as well as added back-to-back dependencies shows that why we choose to implement the destination channel replication for floating point DP4.

```
dp4 (8) r5.0<1>.y:f    r1.0<4;4,1>:f    r4.0<4;4,1>:f
mov (1) r5.1<1>:f      r8.0<1;1,1>:f
dp4 (8) r5.0<1>.z:f    r2.0<4;4,1>:f    r4.0<4;4,1>:f
mov (1) r5.2<1>:f      r8.0<1;1,1>:f
dp4 (8) r5.0<1>.w:f    r3.0<4;4,1>:f    r4.0<4;4,1>:f
mov (1) r5.3<1>:f      r8.0<1;1,1>:f
dp4 (8) r5.0<1>.x:f    r0.0<4;4,1>:f    r4.0<4;4,1>:f
```

## 9.2.3 Null Register as the Destination

Null register can be used as the destination for most of the instructions. Here are some example usages.

- Null as destination for regular ALU instructions: As all ALU instructions can be configured to update the flag registers using the conditional modifiers, it is not necessary to have a destination register if the programmer only cares about the conditionals of the operation. In that case, a null in the destination operand field saves register space as well as one less dependency checking.
- Null as the destination for SEND/STOR instructions: for the send instruction that only send messages out to an external unit and does not require any return data or feedback, a null in the destination register field signifies the case.



## 9.2.4 Use of LINE Instruction

LINE instruction is specifically designed to speed up floating point vector/matrix computation when a program operates in channel serial.

The following example is to compute homogeneous dot product for OpenGL pixel shader running in Channel Serial. In this example, an original OpenGL PS instruction is like

dph R2.x R0 R1

With register remapping, we can store the input coefficient vector R0 in original format in r0, but 8 sets of input coordinate vectors in channel serial format in r2, r3, r4 and r5, and the destination R2.x component in r6.

r0: Cx0 Cy0 Cz0 Co0 DC DC DC DC

r2: X0 X1 ... X7

r3: Y0 Y1 ... Y7

r4: Z0 Z1 ... Z7

r5: W0 W1 ... W7

The objective is to compute the following DPH equations and store the results in r6 as

R6: (X0\*Cx0+Y0\*Cy0+Z0\*Cz0+Co0) ... (X7\*Cx0+Y7\*Cy0+Z7\*Cz0+Co0)

### Example 9-1. Homogeneous Dot Product in Channel Serial

//-----

// Example compute homogeneous dot product in channel serial scenario

//-----

```

line (8) acc:f r0<0;1,0>:f r2<0;8,1>:f // does acc = X# * Cx0 + Co0
mac (8) acc:f r0.1<0;1,0>:f r3<0;8,1>:f // does acc.# = Y# * Cy0 + acc.#
mac (8) r6<1>:f r0.2<0;1,0>:f r4<0;8,1>:f // does r6.# = Z# * Cz0 + acc.#

```





## 9.2.5 Mask for SEND Instruction

Execution mask (upto 16 bits) for the SEND instruction is transferred to the Shared Function.

### 9.2.5.1 Channel Enables for Extended Math Unit

The following example demonstrates how to use the SEND instruction to get service from the Extended Math unit.

Let's consider COS instruction in the following form

```
[(!]p0.{select[any]all})] cos[_sat] dest[.mask], [-]src0[_abs][.swizzle]
```

For a SIMD4x2 VS implementation with the following register mappings

```
p0          →    f0.0
```

```
src0   →    r0
```

```
dest   →    r1
```

The equivalent DevSNB instruction is as the following

```
[(!]f0.0.{select[any4h]all4h})] SEND (8) r1[.mask]:f m0 [-][!(abs)]r0[.swizzle]:f MATHBOX|COS|[SAT]
```

If the source swizzle is replication, the message description field can be modified to MATHBOX|COS|SCALAR to take advantage of the fast mode (scalar mode) supported by the Extended Math. The implied move of the SEND instruction is equivalent to the following instruction:

```
MOV (8) m0[.mask]:f [-][!(abs)]r0.0[.swizzle]:f {NoMask}
```

For a SIMD16 PS implementation, the register mappings are as the followings

```
p0          →    f0...f3                // in order of R, G, B, A
```

```
src0   →    r0,r1; r2,r3; r4,r5; r6,r7
```

```
dest   →    r8,r9; r10,r11; r12,r13; r14,r15
```

There are several ways to translate the instruction, depending on the operand/instruction modifiers present in the instruction. If predicate is not present and the source swizzle is replication, say, src0.y, which is r2-r3, the translation could be as the following instructions

```
send (8) r8:f m0          -(abs)r2:f MATHBOX|COS
```

```
send (8) r9:f m1          -(abs)r3:f MATHBOX|COS {SecHalf}    // use the second half of 8 flag bits
```

```
mov (16) r10:f  r8:f                // All destination color chan's are same
```

```
mov (16) r12:f  r8:f                // MOV is faster than most MathBox func's
```

```
mov (16) r14:f  r8:f                // These MOV's are compressed instructions
```

Notice that instead of issuing Extended Math messages with the same input data, destination color channel replication is performed by the MOV instructions. This is faster for the thread for most cases as



many Extended Math functions consume multiple cycles. This also conserves message bus bandwidth as well as the usage of the shared resource – Extended Math. The destination mask in the instruction indicates which of the r8 to r15 registers are updated. If the source swizzle is not replication, there will be 8 SEND instructions.

With predication on, if the predication modifier is p0.select, translation is to take the selected flag register f#. The other predication modifiers '.any' and '.all' are translated into '.any4v' and '.all4v', respectively. Notice that with predication on, it is not required to run all 4 pixels in a subspan in the same way, so no need to enforce .any4h/.any4v. The following example shows the instruction with predication (but without .select modifier).

```
(f0[.any4v|.all4v]) send (8) r8:f m0      -(abs)r2:f MATHBOX|COS
(f0[.any4v|.all4v]) send (8) r9:f m1      -(abs)r3:f MATHBOX|COS {SecHalf}
(f1[.any4v|.all4v]) mov (16) r10:f r8:f    // All destination color chan's are same
(f2[.any4v|.all4v]) mov (16) r12:f r8:f    // MOV is faster than most MathBox func's
(f3[.any4v|.all4v]) mov (16) r14:f r8:f    // These MOV's are compressed instructions
```

The same instructions works also for predication with select component modifier. We simply replace f0 to f3 above by the selected flag register, say, f1. The modifier of any4h/all4v would also work.

### 9.2.5.2 Channel Enables for Scratch Memory

The following example demonstrates how to use the SEND instruction to get service from the Data Port for scratch memory access.

Let's consider the case when a scratch register is the destination of an instruction.

```
p0          →    f0
src0        →    r0
src1        →    r1
dest        →    s2 / r10
```

We have

```
add (8) m1:f [-][abs]r0[.swizzle]:f [-][abs]r1[.swizzle]:f
mov (8) r3:d r0:d {NoMask}    // move scratch base address to be assembled with offset values
mov (1) r3.0:d 2*32 {NoMask}  // s2 for vertex 0
mov (1) r3.1:d 2*32+16 {NoMask} // s2 for vertex 1
```

```
[[!f0.{sel|any4h|all4h}]] send (8) null[.mask] m0 r3 DATAPORT|RC|WRITE_SIMD2
```

Notice that with a null as the posted destination register, we are able to transfer the [.mask] over the message channel enables. In many cases for scratch memory access, a write-with-commit is required, therefore, the posted destination register could be r10.

Now, let's consider the PS case when a scratch register is the destination of an instruction.



p0 → f0-f4  
src0 → r0-r7  
src1 → r8-r15  
dest → s16-s23 / r16-r23

When predication is not on (or predication with swizzle control on), we have

```
add (16) m4:f [-]([abs])r0/2/4/6_BasedOnSwizzle:f [-]([abs]) r8/10/12/14_BasedOnSwizzle:f
add (16) m6:f [-]([abs])r0/2/4/6_BasedOnSwizzle:f [-]([abs]) r8/10/12/14_BasedOnSwizzle:f
add (16) m8:f [-]([abs])r0/2/4/6_BasedOnSwizzle:f [-]([abs]) r8/10/12/14_BasedOnSwizzle:f
add (16) m10:f [-]([abs])r0/2/4/6_BasedOnSwizzle:f [-]([abs]) r8/10/12/14_BasedOnSwizzle:f
mov (8) r3:d 0x76543210:v {NoMask} // ramp function
mul (16) acc0:d r3:d 16 {NoMask} // ramp function
add (8) acc0:d acc0:d 64 {NoMask,SecHalf} // ramp function
add (16) m2:d acc0:d 2*256 {NoMask} // ramp function
```

```
send (16) null m1 r3 DATAPORT|RC|WRITE_SIMD16
```

As there is no bit left from the unit specified descriptor field, the 4 bit mask must be put into the header field in m1, which requires at least two more instructions.

Alternatively, or for the case that predication without modifier is on, we can do a read-modify-write.

```
mov (8) r3:d 0x76543210:v {NoMask} // ramp function
mul (16) acc0:d r3:d 16 {NoMask} // ramp function
add (8) acc0:d acc0:d 64 {NoMask,SecHalf} // ramp function
add (16) m2:d acc0:d 2*256 {NoMask} // ramp function
```

```
send (16) r16 m1 r3 DATAPORT|RC|READ_SIMD16 // read from scratch
```

// some of the following four instructions may be omitted based on [.mask] field

```
[(!]f0.{sel|any4v|all4v}) add (16) r16:f [-]([abs])r0/2/4/6_BasedOnSwizzle:f [-]([abs])
r8/10/12/14_BasedOnSwizzle:f
```

```
[(!]f0.{sel|any4v|all4v}) add (16) r18:f [-]([abs])r0/2/4/6_BasedOnSwizzle:f [-]([abs])
r8/10/12/14_BasedOnSwizzle:f
```

```
[(!]f0.{sel|any4v|all4v}) add (16) r20:f [-]([abs])r0/2/4/6_BasedOnSwizzle:f [-]([abs])
r8/10/12/14_BasedOnSwizzle:f
```

```
[(!]f0.{sel|any4v|all4v}) add (16) r22:f [-]([abs])r0/2/4/6_BasedOnSwizzle:f [-]([abs])
r8/10/12/14_BasedOnSwizzle:f
```

```
mov (16) m4:f r16:f {NoMask}
```



```
mov (16) m6:f r18:f {NoMask}
mov (16) m8:f r20:f {NoMask}
mov (16) m10:f r22:f {NoMask}

send (16) null m1 null DATAPORT|RC|WRITE_SIMD16 {NoMask} // write back to scratch
```

## 9.2.6 Flow Control Instructions

Unconditional branches are performed through direct manipulation of the 32-bit IP architectural register. For example:

```
mov (1) IP <memory_address> // jump absolute
add (1) IP IP <byte_count> // jump relative
```

Note that jump distances are specified in terms of bytes, as opposed to instruction counts in the case of *break*, *halt*, etc. To minimize confusion, an assembler-only instruction ‘*jmp <inst\_count>*’, where *<inst\_count>* is an immediate term, may be defined which takes an instruction count for a distance. The *jmp* pseudo-opcode can be mapped to an “*add (1) ip ip <inst\_count> \* 16*” instruction.

Also note that IP is always an instruction-sized aligned address (16 bytes), thus the 4 LSB’s are not maintained in the IP architectural register and should not be relied upon by software.

IP, when used as a source operand, reflects the memory address of the instruction in which it is used. The following are examples illustrating the use of IP:

```
add (1) IP 4*16 // jumps to HERE_1
add (1) IP 0x35 // jumps to HERE_1 (4 lsb’s don’t-care)
<instruction>
<instruction>
HERE_1: <instruction>
HERE_2: <instruction>
<instruction>
add (1) IP -2*16 // jumps to HERE_2
...
add (1) IP 0 // infinite loop
add (1) IP 0xF // infinite loop
...
```

**Note for Assembler:** The *if/iff/else/while/break* instructions identify relative addresses as the targets of an implicit jump associated with the instruction. These are optional in the assembly syntax as the jitter can determine the location of the matching instruction (e.g. matching *endif* instruction for a given *if* instruction).



## 9.2.7 Execution Masking

### 9.2.7.1 Branching

#### Example 9-2. If / Else / Endif

```
//-----  
// Example if/else/endif scenario  
//      "if (r5==r4) ...else ... end-if"  
//-----  
  
    ...  
    cmp.e.f0 (8) null r5 r4 // does r5 == r4?  
    (f0) if (8)  HERE_1      // "if" part - save then update IMASK;  
                                // or goto the 'else' if all false  
    ...  
    ...  
HERE_1:                                // now do the 'else' part  
    else (8)  HERE_2        // "else" part - invert IMASK  
                                // or goto the 'endif' if all false  
    ...  
    ...  
HERE_2:  
    endif                                // "end-if" part – restore IMASK  
    ....                               // and continue...
```

If it is known that the code has no nested conditionals, a predicate can be used for a lower overhead, more efficient if/else/endif. (One must consider the probability of all channels taking the same branch, and the number of instructions under the if/else blocks as to which conditional method, predicate or mask, is most efficient).

### 9.2.7.2 Fast-If

Below is an example of a fast-if instruction. For the 'iff' instruction, only and iff-endif construct is allowed, as opposed to a if-else-endif. Note that the target address for branching if all enabled channels fail is one instruction beyond the endif, as the 'iff' does not push and update the IMask unless the branch is taken for at least one execution channel.



### Example 9-3. Fast If

```
//-----  
// Example – Fast If  
//      One instruction overhead conditional  
//-----  
  
    ...  
    cmp.e.f0 (8) null r5 r4      // any flag update  
  
    ...  
(f0)  iff (8) HERE_1           // “fast-if” – only pushes IMask;  
                                           // if execution falls through,  
                                           // else go to HERE_1  
  
    ...  
    ...  
    endif                       // “end-if” part – restores IMask  
  
HERE_1:  
    ...                          // and continue...
```

### 9.2.7.3 Cascade Branching

As there is no ‘elseif’ instruction, a C-like cascade branching such as if / elseif / else / endif, can be realized using the basic building blocks of if / else / endif as shown in the following example. Notice that two ‘endif’s’ are required in order to pop the IStack correctly.

### Example 9-4 If / Elseif / Else / Endif

```
//-----  
// Example if/elseif/else/endif scenario  
//      “if (r5==r4) ...elseif (r6>r7) else ... end-if”  
//-----  
  
    ...  
    cmp.e.f0 (8) null r5 r4  // does r5 == r4?  
  
(f0)  if (8) HERE_1         // “if” part - save then update IMask;  
                                           // or go to the ‘else’ part if all false  
  
    ...
```



```
...
HERE_1:                                     // now do the 'else' part
else (8) HERE_2                             // "else if" part - invert IMask
                                           // or go to the 'else' part if all false
cmp.g.f0 (8) null r6 r7 // is r6 > r7?
(f0) if (8) HERE_3                          // "if" part - save then update IMask;
                                           // or go to the 'else' part if all false
...
...
HERE_3:                                     // now do the 'else' part
else (8) HERE_4                             // "else" part - invert IMask
                                           // or go to the 'end-if' part if all false
...
...
HERE_4:
endif                                       // "end-if" part – restore IMask for elseif
HERE_2:
endif                                       // "end-if" part – restore IMask for if
....
```

#### 9.2.7.4 Compound Branches

Compound branches are supported through the ability logically combine flag registers for each intermediate result.

##### Example 9-5 Compound Branch

```
//-----
// Example: "if (r0 > r1) OR (r2 <= r3)"
//-----
...
cmp.g.f0 (8) null r0:d r1:d // r0 > r1?
cmp.le.f1 (8) null r2:d r3:d // r2 <= r3?
```



```
        or (1) f0:w f0:w f1:w           // combine f0 and f1
(f0) if (8) HERE_1                     // Can now do normal if/else
    ...
    ...
HERE_1:          endif
    ...
```

### Example 9-6. Compound Branch Using 'Any' or 'All'

```
//-----
// Example: assuming we're doing a channel-serial vector in r0-r3
//      We want to know if all components of the vector are > 0x80
//-----

    ...
    cmp.g.f0 (16) null r0 0x80         // r0 > 0x80?
    cmp.g.f1 (16) null r1 0x80         // r1 > 0x80?
    cmp.g.f2 (16) null r2 0x80         // r0 > 0x80?
    cmp.g.f3 (16) null r3 0x80         // r1 > 0x80?
(f0.all4v) if (16) HERE_1
    ...
    ...                               // code executed only for those channels
    ...                               // where per-channel r0,r1,r2,r3 all > 0x80
    ...
HERE_1:          endif
    ...                               // and continue...
```

### 9.2.7.5 Looping

Due to DevSNB's SIMD-16 architecture, it must support the case of up to 16 loops running in parallel. These must be handled as independent loops, each with its own loop-exit condition which could occur after a different number of loop iterations. To account for each channel's progress, a 16b loop-mask 'LMask' is defined with 1b associated to each execution channel. This mask keeps track of which channels remain active inside a loop block.





## Basic Do-While Loop

Example 9-7 illustrates the most basic loop. Two operations must be accomplished before loop entry. (1) Prior to loop entry, there is some subset of enabled channels as dictated by the code sequence prior. In general, the active status of each channel is indicated in the virtual EMask any point in time. These active channels will become the channels over which the loop is run, and LMask must be initialized with the EMask value. (2) Since a given loop may be nested within another loop, the previous LMask & CMask must be saved to the LStack for later restoration upon loop completion. The 'msave' instruction performs both the save and update in a single instruction, and thus all loop-blocks should be fronted with a "msave LStack LMask" and "msave LStack CMask" operation.

Note that the LMask and CMask share the same mask-stack. Thus, CMask must always be a 1's-subset of the LMask for proper stack operation. This is the case if CMask is updated to LMask each pass through the loop (see Example 9-7) and through the 'break' instruction updating both masks.

Each pass through the loop, a loop terminating operation must be evaluated and stored in a flag register. This condition must be evaluated on a channel-by-channel basis as exemplified:

```
cmp.z.f0      (8) null r2 d3      // any operation that updates a flag
```

The result of this operation sets a bit per channel in the specified flag register, which is then used in the 'while' instruction. As loops are performed, channels may become disabled as their termination condition is met.

'While' termination is determined on a channel-by-channel basis by the logical AND of corresponding bit positions of AMask, CMask and the specified flag. If the result is '1' the channel remains enabled for the next pass of the loop; if '0' the channel is disabled until loop fall-through. The 'while' instruction causes the LMask to be updated with the latest result of enabled channels. If any channel remains enabled (LMask != ...000b), an additional pass through the loop is made. Once a channel is terminated for the loop operation, it remains terminated until the loop is complete for all channels.

Upon fall through, the 'while' instruction causes the previously saved LMask & CMask to be popped from the LStack, enabling execution on the same subset of channels enabled prior to loop entry (unless a channel had been otherwise terminate inside the loop via 'halt').

### Example 9-7. Basic Loop Construct

```
//-----  
//      Example: Basic do-while loop structure  
//-----  
  
...  
do                                           // save L/CMask & update  
  
BEGIN_LOOP:  
    mov (1) CMask LMask {NoMask}          // update CMask for this pass  
    ...  
    ...  
    <some flag update>
```



```
(<p> while (8) BEGIN_LOOP           // cond. branch
                                     // + restores LMask on fall-thru
...

```

### Do-While Loop with Break

A loop may also be terminated for any channel via the 'break' instruction. The 'break' instruction causes the corresponding bit positions of enabled channels to be cleared in the LMask. If the updated LMask = ...000b, a branch is made to the specified instruction location. An example is shown below in which the 'break' is at the same conditional-nesting level as the terminating 'while'. Its primary value may simply be to support a "do...break.. while (true)" –type structure for a more direct 1:1 translation from higher-level source code.

#### Example 9-8. Loop Construct With Non-Nested 'Break'

```
//-----
//      Example: While-true loop
//-----
#define BrkCode(i,d)  (i << 16) + d

                                do                // save L/CMask & update
BEGIN_LOOP:
                                mov (1) CMask LMask {NoMask} // update CMask for this pass
                                ...
                                <some flag update>
                                (<p>) break (8) BrkCode(0,HERE_1) // Restores LMask when all
                                                                // channels complete loop.
                                ...
                                ...
                                while (8) BEGIN_LOOP           // while true

HERE_1:
                                ...

```

A break condition may occur from various levels of nested-ifs. This gives rise to the possibility that a the loop may terminate from within nested 'if's, and due to the jump inherent in the 'break' instruction, the associated 'endif's are not encountered to clean-up the IStack as nesting levels are exited.



### Example 9-9 Loop Construct With 'Break' From Within Nested If's

```
//-----  
//      Example: General Loop Structure w/ break inside if's  
//-----  
#define BrkCode(i,d)  (i << 16) + d  
  
do                                     // save L/CMask & update  
BEGIN_LOOP:  
    mov (1) CMask LMask {NoMask} // update CMask for this pass  
    ...  
    if ...  
    if ...  
    if ...  
    ...  
    (<p>) break (8) BrkCode(3,HERE_1) // we're 3 levels deep, so  
    ...  
    endif  
    endif  
    endif  
    ...  
    (<p>) break (8) BrkCode(0,HERE_1)  
    ...  
    while (8) <flag_spec> BEGIN_LOOP // cond. branch  
                                        // + restores C/LMask on fall-thru  
  
HERE_1:
```



## Do-While Loop with Continue

A continue instruction 'cont' is provided skip to the next iteration of the loop. Because not all channels participating in the loop may be enabled at the time this instruction is executed, some channels may require continuation of the loop. A special mask 'CMask' is defined which accounts for channels temporarily disabled for the current loop pass.

Since loops may be nested, the CMask must be saved and restored around a loop similar to LMask. Since the CMask value within a properly constructed loop is always a subset of the LMask, it can share the LStack for storage, so long as it is pushed after LMask as shown in Example 9-10. This save/restore operations are not required if the loop being entered does not have any occurrence of a continue instruction.

### Example 9-10. Do-While with Continue

```
//-----  
//      Example: General Loop Structure w/ basic break and cont.  
//-----  
#define ContCode(i,d)  (i << 16) + d  
  
do                                // save L/CMask & update  
BEGIN_LOOP:  
    mov (1) CMask EMask          // re-initialize CMask for this pass  
    ...  
    ...  
    (<p>) cont (8) ContCode(0,HERE_1)  
    ...  
HERE_1:  
    (<p>) while (8) BEGIN_LOOP    // cond. branch  
    // + restores C/LMask on fall-thru  
    ...
```



## 9.2.7.6 Indexed Jump

### Example 9-11. Indexed Jump

```
//-----  
// Code example shows the use of jmp to perform a case statement  
// of any number of options in 3 jumps  
//-----  
.default_execution_size 8  
  
...  
jmp r0<0,1,0>           // jump relative, based on r0.a.x  
                        // ----- Jump Table -----  
  
jmp HERE_0             // redirect for case 0  
jmp HERE_1             // redirect for case 1  
jmp HERE_2             // redirect for case 2  
jmp HERE_3             // redirect for case 3  
  
...  
HERE_0:                // ... case 0 ...  
  
...  
    jmp DONE  
HERE_1:                // ... case 1 ...  
  
...  
    jmp DONE  
HERE_2:                // ... case 2 ...  
  
...  
    jmp DONE  
HERE_3:                // ... case 3 ...  
  
...  
DONE:  
  
...                    // and continue...
```



## *Revision History*

<b>Revision Number</b>	<b>Description</b>	<b>Revision Date</b>
1.1	Edited extraneous content	May 2011
1.0	First 2011 OpenSource edition	May 2011

§§