

Intel[®] HD Graphics OpenSource PRM

Volume 2 Part 2: 3D/Media - Media

**For the all new 2010 Intel Core Processor Family
Programmer's Reference Manual (PRM)**

February 2010

Revision 1.0



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.

You are not obligated to provide Intel with comments or suggestions regarding this document. However, should you provide Intel with comments or suggestions for the modification, correction, improvement, or enhancement of: 9a) this document; or (b) Intel products, which may embody this document, you grant to Intel a non-exclusive, irrevocable, worldwide, royalty-free license, with the right to sublicense Intel's licensees and customers, under Recipient intellectual property rights, to use and disclose such comments and suggestions in any manner Intel chooses and to display, perform, copy, make, have made, use, sell, and otherwise dispose of Intel's and its sublicensee's products embodying such comments and suggestions in any manner and via any media Intel chooses, without reference to the source.

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 Sandy Bridge 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 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 © 2010, Intel Corporation. All rights reserved.



Contents

| | |
|---|------------|
| 1. Media and General Purpose Pipeline | 6 |
| 1.1.1 Hardware Feature Map in Products | 6 |
| 1.2 Media Pipeline Overview | 6 |
| 1.3 Programming Media Pipeline | 8 |
| 1.3.1 Command Sequence | 8 |
| 1.3.2 Interrupt Latency | 13 |
| 1.4 Video Front End Unit | 13 |
| 1.4.1 Interfaces | 14 |
| 1.4.2 Mode of Operations | 16 |
| 1.4.3 Scoreboard Control [DevILK+] | 35 |
| 1.5 Thread Spawner Unit | 47 |
| 1.5.1 Basic Functions | 48 |
| 1.5.2 Interfaces | 55 |
| 1.6 Media State ([Pre-DevSNB]) | 56 |
| 1.6.1 Media State Model | 56 |
| 1.6.2 VFE_STATE | 58 |
| 1.6.3 VLD_STATE | 59 |
| 1.6.4 VFE_STATE_EX [DevCTG+] | 62 |
| 1.6.5 INTERFACE_DESCRIPTOR | 65 |
| 1.7 Media State and Primitive Commands | 68 |
| 1.7.1 MEDIA_STATE_POINTERS Command ([Pre-DevSNB]) | 68 |
| 1.7.2 MEDIA_OBJECT Command ([Pre-DevGT]) | 69 |
| 1.7.3 MEDIA_OBJECT_EX Command [DevCTG-DevILK] | 77 |
| 1.7.4 MEDIA_OBJECT_PRT Command [DevCTG+] | 123 |
| 1.7.5 [GPGPU_WALKER Command | 125 |
| 1.8 Media Messages | 127 |
| 1.8.1 Thread Payload Messages | 127 |
| 1.8.2 Thread Spawn Message | 151 |
| 2. AVC Bit-Serial Decoder [DevCTG/DevILK] | 156 |
| 2.1 Design Assumptions | 156 |
| 2.2 Bit-Serial Decoding (BSD) Unit | 156 |
| 2.3 H/W Asynchronous Reset | 158 |
| 2.4 H/W Pipeline Flush | 158 |
| 2.5 MMIO Interface | 158 |
| 2.6 Programming the BSD Unit | 159 |
| 2.6.1 Example Command Sequence for Decoding a Single Video Stream | 159 |
| 2.7 AVC_BSD Commands | 161 |
| 2.7.1 BSD_IND_OBJ_BASE_ADDR Command | 161 |
| 2.7.2 AVC_BSD_STATE Commands | 162 |
| 2.7.3 AVC_BSD_OBJECT Command | 181 |
| 2.8 AVC_BSD Output Definitions | 190 |
| 2.8.1 AVC-IT Command Buffer Description | 191 |
| 2.8.2 AVC-IT Data Buffer Description | 191 |
| 2.8.3 AVC-ILDB Data Buffer Description | 216 |
| 2.8.4 Special Scan order for MB/MB Pair Kernel Processing | 216 |
| 2.9 Error Handling and Error Concealment | 217 |
| 2.9.1 Design Assumptions | 217 |



| | | |
|-----------|---|------------|
| 2.9.2 | Error Concealment and Recovery Strategies | 218 |
| 2.9.3 | Error Handling..... | 218 |
| 2.9.4 | Error Handling for Non-Existing Reference Picture..... | 221 |
| 2.10 | Concurrent, Multiple Video Stream Decoding Support..... | 221 |
| 2.11 | Performance and Latency Estimation..... | 222 |
| 2.12 | Media Messages..... | 222 |
| 2.13 | Media Applications with Specific Hardware Support | 222 |
| 2.13.1 | AVC – Off-host CABAC/CAVLD Acceleration | 222 |
| 3. | VC1 Bit-Serial Decoder ([DevCTG, DevILK]) | 223 |
| 3.1 | Introduction..... | 223 |
| 3.2 | VC1_BSD Commands..... | 223 |
| 3.2.1 | VC1_IND_OBJ_BASE_ADDR Command..... | 223 |
| 3.2.2 | VC1_BSD_STATE Commands | 224 |
| 3.2.3 | VC1_BSD_OBJECT Command(Pre-DevILK) | 239 |
| 3.2.4 | VC1_BSD_OBJECT Command(DevILK only) | 241 |
| 3.3 | Context Switch..... | 243 |
| 3.4 | Error Handling and Error Concealment..... | 243 |
| 3.4.1 | Design Assumptions..... | 244 |
| 3.4.2 | Error Concealment and Recovery Strategies..... | 245 |
| 3.4.3 | Error Handling..... | 245 |
| 3.5 | AVC Decoder Commands | 248 |
| 3.5.1 | MFD_AVC_BSD_OBJECT Command | 248 |
| 3.6 | AVC Encoder PAK Commands | 252 |
| 3.6.1 | MFC_AVC_FQM_STATE Command | 253 |
| 3.6.2 | MFC_AVC_PAK_INSERT_OBJECT Command | 254 |
| 3.6.3 | MFC_STITCH_OBJECT Command..... | 257 |
| 3.6.4 | MFC_AVC_PAK_OBJECT Command | 260 |
| 3.7 | VC1 Common Commands..... | 303 |
| 3.8 | VC1 Decoder Commands..... | 303 |
| 3.8.1 | MFD_VC1_BSD_OBJECT Command..... | 303 |
| 3.9 | MPEG2 Common Commands..... | 305 |
| 3.10 | MPEG2 Decoder Commands | 305 |
| 3.10.1 | MFD_MPEG2_BSD_OBJECT Command (pipeline) | 305 |
| 3.11 | MFD IT Mode Decode Commands | 308 |
| 3.11.1 | MFD_IT_OBJECT Command..... | 308 |
| 3.12 | AVC Encoder MBAFF Support | 326 |
| 3.13 | Decoder StreamOut Mode Data Structure Definition | 327 |
| 3.14 | Encoder StreamOut Mode Data Structure Definition..... | 336 |
| 3.14.1 | PAK Multi-Pass..... | 338 |
| 3.14.2 | Driver Usage..... | 339 |
| 3.15 | Monochrome Picture Processing..... | 340 |
| 3.16 | Context Switch | 340 |
| 3.17 | Pipeline Flush | 341 |
| 3.18 | AVC Error Handling and Error Concealment..... | 341 |
| 3.18.1 | AVC Design Assumptions | 342 |
| 3.18.2 | AVC Error Concealment and Recovery Strategies | 342 |
| 3.18.3 | AVC Error Handling | 343 |
| 3.19 | VC1 Error Handling and Error Concealment | 345 |
| 3.20 | MPEG2 Error Handling and Error Concealment..... | 346 |
| 3.21 | Concurrent, Multiple Video Stream Decoding Support..... | 346 |



Revision History

| Document Number | Revision Number | Description | Revision Date |
|------------------------|------------------------|--------------------|----------------------|
| IHD-OS-022810-R1V2PT2 | 1.0 | First Release. | February 2010 |



1. Media and General Purpose Pipeline

1.1.1 Hardware Feature Map in Products

The following table lists the hardware features in the media pipe.

Video Front End Features in Device Hardware

| Features/ Device | [DevBW] | [DevCL] | [DevCTG] | [DevILK] |
|--|---------|---------|----------|----------|
| Generic Mode | Y | Y | Y | Y |
| Root Threads | Y | Y | Y | Y |
| Parent/Child Threads | Y | Y | Y | Y |
| SRT (Synchronized Root Threads) | Y | Y | Y | Y |
| PRT (Persistent Root Thread) | N | N | Y | Y |
| Interface Descriptor Remapping | N | N | Y | Y |
| Interface Descriptor Remapping | N | N | Y | Y |
| IS Mode (HW Inverse Scan) | Y | Y | Y | Y |
| VLD Mode (HW MPEG2 VLD) | N | Y | Y | Y |
| AVC MC Mode | N | N | Y | Y |
| AVC IT Mode (HW AVC IT) | N | N | Y | Y |
| AVC ILDB Filter (in Data Port) | N | N | Y | Y |
| VC1 MC Mode | N | N | Y | Y |
| VC1 IT Mode (HW VC1 IT) | N | N | Y | Y |
| Stalling HW Scoreboard | N | N | N | Y |
| Non-stalling HW Scoreboard | N | N | N | N |
| HW Walker | N | N | N | N |
| HW Timer | N | N | N | N |
| Pipelined State Flush | N | N | N | N |
| HW Barrier | N | N | N | N |

1.2 Media Pipeline Overview

The media (general purpose) pipeline consists of two fixed function units: Video Front End (VFE) unit and Thread Spawner (TS) unit. VFE unit interfaces with the Command Streamer (CS), writes thread payload data into the Unified



Return Buffer (URB) and prepares threads to be dispatched through TS unit. VFE unit also contains a hardware Variable Length Decode (VLD) engine for MPEG-2 video decode. TS unit is the only unit of the media pipeline that interfaces to the Thread Dispatcher (TD) unit for new thread generation. It is responsible of spawning root threads (short for the root-node parent threads) originated from VFE unit and spawning child threads (can be either a leaf-node child thread or a branch-node parent thread) originated from the Execution Units (EU) by a parent thread (can be a root-node or a branch-node parent thread).

The fixed functions, VFE and TS, in the media pipeline, in most cases, share the same basic building blocks as the fixed functions in the 3D pipeline. However, there are some unique features in media fixed functions as highlighted by the followings.

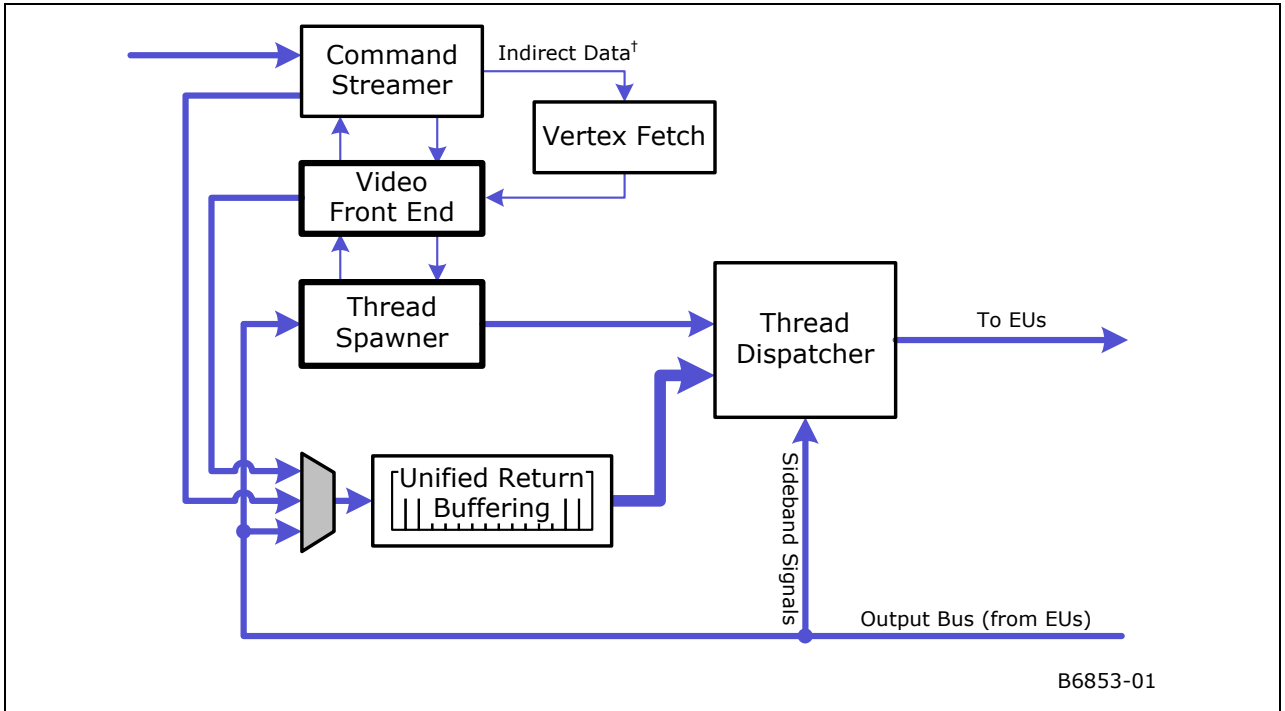
- VFE manages URB and only has write access to URB; TS does not interface to URB.
- When URB Constant Buffer is enabled, VFE forwards TS the URB Handler for the URB Constant Buffer received from CS.
- TS interfaces to TD; VFE does not.
- TS can have a message directed to it like other shared functions (and thus TS has a shared function ID), and it does not snoop the Output Bus as some other fixed functions in the 3D pipeline do.
- A root thread generated by the media pipeline can only have up to one URB return handle.
- If a root thread has a URB return handle, VFE creates the URB handle for the payload to initiating the root thread and also passes it alone to the root thread as the return handle. The root thread then uses the same URB handle for child thread generation.
- If URB Constant Buffer is enabled and an interface descriptor indicates that it is also used for the kernel, TS requests TD to load constant data directly to the thread's register space. For root thread, constant data are loaded after R0 and before the data from the other URB handle. For child thread, as the R0 header is provided by the parent thread, Thread Spawner splits the URB handles from the parent thread into two and inserts the constant data after the R0 header.
- A root thread must terminate with a message to TS. A child thread should also terminate with a message to TS.
- High streaming performance of indirect media object load is achieved by utilizing the large vertex cache available in the Vertex Fetch unit (of the 3D pipeline).

[DevBW] Erratum: DevBW does not have MPEG-2 VLD hardware. Therefore, software cannot use the VLD mode of the Media_Object command.

[DevBW-A] Erratum: Using vertex cache in Vertex Fetch unit to speed up streaming of indirect media data load is not available on DevBW-A. On DevBW-A, indirect media data are loaded directly from CS to VFE.



Figure 1-1. Top level block diagram of the Media Pipeline



1.3 Programming Media Pipeline

1.3.1 Command Sequence

Media pipeline uses a simple programming model. Unlike the 3D pipeline, it does not support pipelined state changes. Any state change requires an MI_FLUSH or PIPE_CONTROL command. When programming the media pipeline, it should be cautious to not use the pipelining capability of the commands described in the Graphics Processing Engine chapter.

To emphasize the non-pipeline nature of the media pipeline programming model, the programmer should note that if any one command is issued in the “Primitive Command” step, none of the state commands described in the previous steps cannot be issued without preceding with a MI_FLUSH or PIPE_CONTROL command.

The basic steps in programming the media pipeline are listed below. Some of the steps are optional; however, the order must be followed strictly. Some usage restrictions are highlighted for illustration purpose. For details, reader should refer to the respective chapters for these commands.



1.3.1.1 [Pre-DevSNB]

- Special Requirements for Each Context Initialization
 - Always initialize the URB fence (with a URB_FENCE command) before the first pipeline select command (PIPELINE_SELECT).
 - Always initialize the pipeline state pointer (with a STATE_BASE_ADDRESS command) before the first pipeline select command.
- Step 1: MI_FLUSH/PIPE_CONTROL
 - This step is mandatory.
 - Programmer may choose not to flush certain caches to improve performance.
 - Multiple such commands in step 1 are allowed, but not recommended for performance reason.
 - [DevBW-B, DevBW-C, DevCLN]
 - MI_LOAD_REGISTER_IMM
 - It is used to load an MMIO register to disable the vertex cache for indirect media object load. The register is 0x2124 and the bit is 15.
 - Address = 0x2124
 - Data = 0x10000000
 - MI_FLUSH
 - MI_LOAD_REGISTER_IMM
 - This command is optional for this step. It is only required when indirect object load is used subsequently by MEDIA_OBJECT commands.
 - If present, it is used to load an MMIO register to enable the vertex cache for indirect media object load. The register is 0x2124 and the bit is 15.
 - Address = 0x2124
 - Data = 0x10001000
- Step 1.5: State commands SF_STATE + URB_FENCE [Errata: Pre-DevILK]
 - When switching from 3D context to Media context, the following sequence must be sent before the PIPE_SELECT command.
 - SF_STATE command must be sent down to set the “Number of URB Entries” to “0”.
 - URB_FENCE command must be sent down to set the “URB Fence” for all 3D units to “0”, including CS, GS, CL, and SF.
- Step 2: State command PIPELINE_SELECT
 - This step is optional. This command can be omitted if it is known that within the same context media pipeline was selected before Step 1.
 - Multiple such commands in step 2 are allowed, but not recommended for performance reason.
 - If this command is issued, it must be followed by a URB_FENCE command (step 3).



- Step 3: State command URB_FENCE
 - This step is optional. This command can be omitted if URB fence needs not to be changed. However, as mentioned above, if a PIPELINE_SELECT command is issued, this command is then required.
 - If present, only one URB_FENCE command in step 3 is allowed. Hardware behavior is undefined if more than one URB_FENCE commands are issued in this step.
- Step 4: State commands configuring pipeline states
 - STATE_BASE_ADDRESS
 - This command is mandatory for this step (i.e. at least one).
 - Multiple such commands in this step are allowed. The last one overwrites previous ones.
 - This command must precede any other state commands below.
 - Particularly, the fields **Indirect Object Base Address** and **Indirect Object Access Upper Bound** are used to control indirect object load.
 - *Note: This command may be inserted before (and after) any commands listed in the previous steps (Step 1 to 3). For example, this command may be placed in the ring buffer while the others are put in a batch buffer.*
 - The following state commands can be issued in arbitrary order.
 - MEDIA_STATE_POINTERS
 - This command is mandatory for this step (i.e. at least one).
 - Multiple such commands in this step are allowed. The last one overwrites previous ones.
 - CS_URB_STATE
 - This command is optional for this step. Note that if CS_URB_STATE command is present, there will be at least one MEDIA_STATE_POINTERS command in this step (as mentioned above).
 - Multiple such commands in this step are allowed. The last one overwrites previous ones.
 - If present, “Number of URB Entries” must be 0 if no URB entry is allowed to CS by URB_FENCE command.
 - “Number of URB Entries” must be set to 1 as media pipeline does not support pipelined CONSTANT_BUFFER command (see step 5).
 - STATE_PREFETCH
 - This command is optional for this step.
 - STATE_SIP
 - This command is optional for this step. It is only required when SIP is used by the kernels.



- 3DSTATE_VERTEX_ELEMENTS ([**DevBW-B**, **DevBW-C**, **DevCLN**] only. For other products, this command cannot be issued as indirect object load is fully described by each MEDIA_OBJECT command.)
 - This command is optional for this step. It is only required when indirect object load is used subsequently by MEDIA_OBJECT commands.
 - If present, only the following programming is allowed. Hardware behavior with other programming is undefined,
 - Two elements need to be programmed
 - Vertex Element 0
 - Vertex Buffer Index = 0
 - Valid = True
 - Surface Format = 0x002
 - Source Element Offset = 0x0
 - Component Control 0,1,2,3 = 0x1
 - Destination Offset = 0x0
 - Vertex Element 1
 - Vertex Buffer Index = 0
 - Valid = True
 - Surface Format = 0x002
 - Source Element Offset = 0x10
 - Component Control 0,1,2,3 = 0x1
 - Destination Offset = 0x10
- 3DSTATE_VERTEX_BUFFERS ([**DevBW-B**, **DevBW-C**, **DevCLN**] only. For other products, this command cannot be issued as indirect object load is fully described by each MEDIA_OBJECT command.)
 - This command is optional for this step. It is only required when indirect object load is used subsequently by MEDIA_OBJECT commands.
 - If present, only the following programming is allowed. Hardware behavior with other programming is undefined,
 - Only 1 vertex buffer
 - Buffer Access Type : Vertex Data
 - Buffer Pitch : 0x20
 - Buffer Start Address : <Indirect Data Address>
 - When VFE is in Generic Mode, the vertex buffer base address can be byte aligned. The restriction is that indirect data size of each MEDIA_OBJECT command must be a multiple of 32 bytes.
 - When VFE is either in VLD mode or IS mode, the indirect data size may not be multiple of 32 bytes (that's OK). However, it is required that the vertex buffer to be programmed to be 32-byte aligned. All indirect data must be included in the vertex buffer programmed.
 - Max index is always set to 0 (ie disabled)



- Step 5: State command CONSTANT_BUFFER
 - This step is optional. However, it is required (as a software workaround) when 3DPRIMITIVE commands are used subsequently to load indirect object data.
 - If present, only one such command is allowed. Hardware behavior is undefined if more than one CONSTANT_BUFFER commands are issued in the program sequence without a FLUSH in between.
- Step 6: Primitive commands
 - 3DPRIMITIVE ([DevBW-B, DevBW-C, DevCLN] only. For other products, this command cannot be issued as indirect object load is fully described by each MEDIA_OBJECT command.)
 - This command is optional for this step. It is only required when indirect object load is used subsequently by MEDIA_OBJECT commands.
 - If present, this command must precede one or many MEDIA_OBJECT commands. If more than one MEDIA_OBJECT commands are followed, the indirect object data for these commands must be stored in memory contiguously (with certain 32-byte aligned overlaps allowed, see XXX for details).
 - If present, only the following programming is allowed. Hardware behavior with other programming is undefined,
 - Sequential access for the Vertex Buffer
 - Primitive topology type is 0x01h = PointList
 - Vertex Count per instance = Size of the block to transfer for the media indirect command in 32 byte quantities.
 - Start Vertex Location = 0
 - Instance Count = 1
 - Start instance Location = 0
 - Base Vertex Location = 0
 - MEDIA_OBJECT
 - This step is optional, but it doesn't make practical sense not issuing media primitive commands after being through previous steps to set up the media pipeline.
 - Multiple such commands in step 6 can be issued to continue processing media primitives.

rogramming Notes on Improving Indirect Media Object Load Performance [DevBW-C, DevCL]: The large vertex cache is used to stream indirect media object loads for one or many MEDIA_OBJECT commands. By grouping multiple such commands together significant streaming performance can be achieved. Here is an example.

- 1 Vertex Buffer programmed with 2 Vertex components
- Vertex format is fixed to A32R32B32G32_UINT, this format is left untouched by the vertex fetch

Here the number of vertices equal to the total size to be transferred for MEDIA_OBJECT commands in 32-byte chunks. If the first MEDIA_OBJECT command transfers indirect data size of 4 64 byte quantities, the number of vertices would be 8, If the second MEDIA_OBJECT command to be transferred and the total size is 8 64 byte quantities, number of vertices is 16.



However, using the Vertex Buffer in sequential mode as described above does post a restriction that data for multiple MEDIA_OBJECT commands sharing the same 3DPRIMITIVE command must be stored in memory sequentially. When data are not stored sequentially in memory, there are several approaches as listed below. Certain experiments may be required in order to find which approach provides the best performance for a given application.

- Preceding each MEDIA_OBJECT command with one 3DPRIMITIVE command.
- Grouping several 3DPRIMITIVE commands together followed by the MEDIA_OBJECT commands using the fetched data from the 3DPRIMITIVE commands.
- Using indexed vertex buffer to gather indirect media object data from non-contiguous memory locations.

As a side effect, when vertex cache is used for media indirect object load, the statistical counters in the VF unit may be affected during media operations. When 3D operations and media operations are from different contexts, this side effect is not an issue as the statistical counters are context save/restored. However, if 3D and media operations are mixed within one context, it is advisable to turn off the statistical counters before entering media operation (using vertex cache for indirect object load) and turn them back on before returning to 3D operations. This can be achieved using the 3DSTATE_VF_STATISTICS command.

1.3.2 Interrupt Latency

Command Streamer is capable of context switching between primitive commands.

For all independent threads, it is not much a problem. The interrupt latency is dictated by the longest command that is likely to have the largest number of threads. For VLD mode, such a command may be corresponding to a largest slice in a high definition video frame. This is application dependent, there are not much host software can do. For Generic mode, programmer should consider to constrain the compute workload size of each thread.

In modes with child threads, a root thread may be persist in the system for long period of time – staying until its child threads are all created and terminated. Therefore, the corresponding primitive command may also last for long time. Software designer should partition the workload to restrict the duration of each root thread. For example, this may be achieved by partitioning a video frame and assigning separate primitive commands for different data partitions.

In modes with synchronized root threads, a synchronized root thread is dependent on a previous root or child thread. This means context switch is not allowed between the primitive command for the synchronized root thread and the one for the depending thread. So no command queue arbitration should be allowed between them. Software designer should also restrict the duration of such non-interruptible primitive command segments.

1.4 Video Front End Unit

The Video Front End unit is the first fixed function unit in the media pipeline. It processes MEDIA_OBJECT commands to generate root threads by preparing the control (including interface descriptor pointers) and payload (data pushed into the GRF) for the root threads.

VFE supports three modes of operation: Generic mode, Inverse Scan mode and VLD mode.

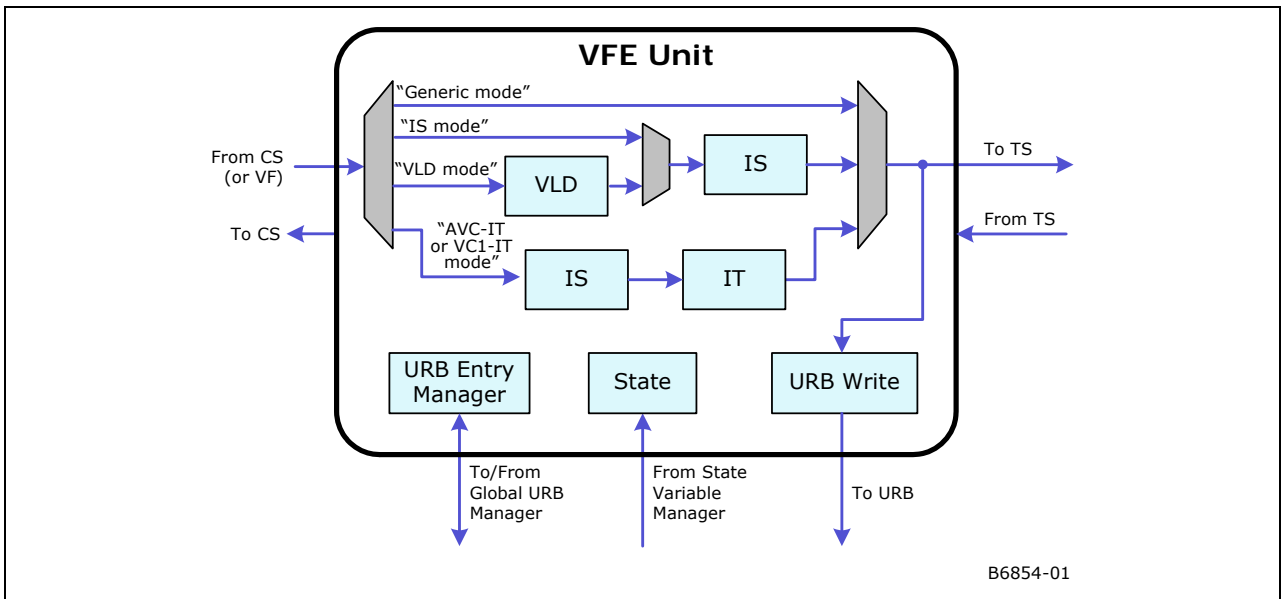
- **Generic mode [Pre-DevSNB]:** In the Generic mode, VFE serves as a conduit for general-purpose kernels fully configured by the host software. There is no application specific hardware enabled in this mode.



- **IS (Inverse Scan) mode:** The IS mode is a special mode for video decoding when off-host IDCT acceleration is supported by kernels running on GENx execution units.
- **VLD mode [Pre-DevSNB]:** It is a special mode for video decoding when MPEG-2 off-host VLD acceleration is supported by GENx hardware.
- **[DevCTG, DevILK], AVC-IT (AVC Inverse Integer Transform) mode:** The AVC-IT mode is a special mode for AVC video decoding when off-host IDCT acceleration is supported by VFE hardware and MC and Loop Filter are supported by kernels running on GENx execution units.
- **[DevCTG, DevILK] AVC-MC (AVC Motion Compensation) mode:** The AVC-MC mode is a special mode for AVC video decoding with host-based IDCT and MC and Loop Filter are supported by kernels running on GENx execution units.
- **[DevCTG, DevILK] VC1-IT (VC1 Inverse Integer Transform) mode:** The VC1-IT mode is a special mode for VC1 video decoding when off-host IDCT acceleration is supported by VFE hardware and MC and Loop Filter are supported by kernels running on GENx execution units.

The following figure illustrates the three modes of operation. The details can be found in the rest of the sections.

VFE Functional Blocks and Modes of Operations



MEDIA_STATE_POINTERS command configures VFE in one of the three modes using. Mode switching requires media pipeline state change.

1.4.1 Interfaces

VFE unit acquires its states from State Variable Manager, accesses URB handles from the Global URB Manager, receives state and primitive commands from CS unit, writes thread payloads to URB, and sends new thread to TS unit. It does not directly interface to Thread Dispatcher. When VFE is ready for a thread, it sends the interface descriptor pointer for the thread to TS.



1.4.1.1 Interface to Command Streamer

VFE interfaces to CS to acquire the control data, inline data and indirect data of MEDIA_OBJECT commands. The interface supports the throughput of a given mode of operation of VFE. For example, in VLD mode and IS mode, VFE consumes one dword at a time, one dword to the variable length decoder or one dword to the inverse-scan operator. In Generic mode, VFE is capable of a much higher throughput to push indirect data (as thread payload data) into URB. As throughput for indirect data is much higher than that of inline data, when large amount of user data need to be passed through VFE unit, if applicable, it is encouraged to use indirect object load.

1.4.1.2 Interface to Thread Spawner

When a new root thread is fully assembled by VFE, VFE passes to TS the interface descriptor pointer, the URB handle information, etc. In response to this, TS processes the thread information and sends a thread request to TD.

VFE also transmits scratch memory base address received from State Variable Manager to TS, and passes on the Constant URB handle received from CS.

VFE receives URB handle dereference signal from TS.

1.4.1.3 Interface to State Variable Manager

State Variable Manager is responsible of fetching media state structure from memory. VFE only acquires its state variable upon the first primitive command. Therefore, host software is allowed to change media states before issuing primitive commands. As media pipeline does not support pipelined state change, a pipeline flush is required before any state change to make sure that there are no outstanding primitive commands in the pipeline.

1.4.1.4 Interface to Global URB Manager

VFE is responsible for managing URB handles for all root threads. Upon state change, VFE allocates URB handles through the Global URB Manager. VFE manages the URB handles in a circular buffer. URB handle referencing is in a strict order (taking from the head of the circular buffer), even though the handle dereferencing may occur out of order.

When starting a root thread, VFE reference one and only one URB handle, forwarding it to TS. TS then forwards this handle to TD for thread dispatching.

The URB handle for a root thread is used in two ways: (1) serving as buffer space for VFE to assemble thread payload, and (2) serving as the return URB buffer for the root thread to assemble child threads and their payload.

TS sends an indication to VFE when it is safe to dereference the URB handle, and VFE dereferences it. After a URB handle has been dereferenced, VFE can assign it to a new thread.

1.4.1.5 Interface to URB

VFE sends the assembled root thread payload to URB via a wide data bus. In Generic mode, the data comes from the command as inline or indirect data objects. In IS mode, the inline data is directly assembled as URB register wide payloads, and the indirect data are assembled through the Inverse Scan logic. In VLD mode, the data is decoded from the indirect object (i.e. bitstream data).



1.4.2 Mode of Operations

1.4.2.1 Generic Mode

In the Generic mode, VFE serves as a conduit for general-purpose kernels fully configured by the host software. As there is no special fixed function logic used, the Generic mode can also be viewed as a ‘pass-through’ mode. In this mode, VFE generates a new thread for each MEDIA_OBJECT command. The payload contained in the MEDIA_OBJECT command (inline and/or indirect) is streamed into URB. The interface descriptor pointer is computed by VFE based on the interface descriptor offset value and the interface descriptor base pointer stored in the VFE state. VFE then forwards the interface descriptor pointer and the URB handle to TS to generate a new root thread. Many media processing applications can be supported using the Generic mode: MPEG-2 HWMC, frame rate conversion, advanced deinterface filter, to name a few.

1.4.2.1.1 Interface Descriptor Selection

After populating the URB with the data, VFE notifies TS to initiate the thread. TS needs an interface descriptor pointer to fetch the information for thread initiation. A list of interface descriptors is arranged by the host software as a descriptor array in memory, as shown in the media state model in **Error! Reference source not found.**

VFE obtains the interface descriptor base pointer from the VFE state structure. The offset into the list of interface descriptors comes from MEDIA_OBJECT command. Each interface descriptor has a fixed size. VFE uses a multiple of the fixed size and the offset to add to the base pointer, and creates the final interface descriptor pointer to be sent to TS.

TS fetches the interface descriptor through the Instruction State Cache (ISC) using the interface descriptor pointer. TS then initializes the thread through the Thread Dispatcher. The interface descriptor pointer is given to TS by VFE for a root thread and by a thread for a child thread. The R0 header is formed by TS for a root thread and is stored in URB by the parent thread for a child thread.

1.4.2.1.2 Scratch Space Allocation

TS handles the allocation of scratch space. Since TS does not have a normal state interface, VFE receives the scratch space configuration with the VFE state, then forwards the configuration to TS with the interface descriptor pointer.

1.4.2.2 IS Mode [Pre-DevSNB]

In Inverse Scan (IS) mode, the Inverse Scan unit is used. In particular, GENx architecture can be used to support off-host IDCT acceleration for MPEG-2.

In this mode, a new thread is generated for each MEDIA_OBJECT command. One MEDIA_OBJECT command corresponds to a macroblock. The indirect payload in the command contains the non-zero DCT coefficients for all coded blocks in the macroblock. Detailed data format can be found in section 1.7.2.2.

The indirect payload is streamed into the IS unit. IS unit process the non-zero DCT coefficients on a block by block basis. The 16-bit non-zero coefficients are placed in its location within an 8x8 block according to the (x,y) addresses. Hardware fills the rest of the coefficients in the block to zero. The assembled DCT data blocks are then written into URB.

Note that the index for a non-zero coefficient is in row-major order (x, y) address. Host software is responsible of converting the coding scan order to this unified row-major order (e.g. zig-zag scan or alternative scan order such as vertical scan found in MPEG-2 or other coding standard).



Blocks that are not coded will not have coefficient data in the message to the kernel, and the coded blocks are packed back to back. As the message size is variable, VFE calculates the final message size according to the coded block pattern field before sending it to TS.

Interface descriptor select and scratch memory allocation are handled in the same way as in Generic mode.

1.4.2.3 VLD Mode [Pre-DevSNB]

In VLD mode, both the VLD unit and the IS unit in VFE are used. The VLD unit is specifically designed to support MPEG-2 variable length decoding; it does not support other standards such as WMV. Each MEDIA_OBJECT command contains compressed bitstream data associated with a slice. A slice, according to MPEG-2 compressed video bitstream syntax show in **Error! Reference source not found.**, is the smallest unit that marked by byte-aligned start-code, allowing easy parsing by host software. A slice contains one or more macroblocks. Unlike the other two modes, one or many threads may be generated for each MEDIA_OBJECT command. Each thread corresponds to a macroblock. The indirect payload in the MEDIA_OBJECT command contains the bitstream data for a slice. The indirect payload is streamed into the VLD unit. The decoded non-zero coefficients are then sent to the IS unit. And then the IDCT data blocks (8x8 size) output from the IS unit are then written into URB. For each macroblock, VFE generates the interface descriptor pointer based the decoded macroblock type.

VFE partially decodes (VLD and IS) the MPEG-2 bitstream for a slice and assembles resulting data on a macroblock by macroblock basis for threads running on EUs to complete the rest of the work. The macroblock-based thread (referred to as a post-VLD thread hereafter) performs inverse quantization, inverse DCT, and motion compensation in order to generate the final output picture. VFE also handles skipped macroblocks so that each post-VLD thread operates on one and only one macroblock.

For bitstreams that are MPEG-2 standard compliant, the output from the VFE fixed function hardware is bit accurate. Bit precision difference may be caused by the IDCT implementation in the kernel. The IDCT kernel must meet the IEEE standard requirement for IDCT.

For bitstreams that are not MPEG-2 standard compliant due to, for example, data corruption, output from VFE fixed function may be unpredictable. That may result in data corruption in the destination buffer after kernel operation. However, VFE fixed function will continue functioning (without hanging).

VFE decodes the slice through the following three major stages: variable length decode, inverse scan and output formatting.

1.4.2.3.1 Variable Length Decode

Variable Length Decode (VLD) stage contains the following sub-stages: data parser, symbol decoder, and motion vector (MV) predictor.

Data Parser

Slice data are processed a dword at a time. Using the byte offset and bit offset provided by the MEDIA_OBJECT command, data parser determines the start bit and sends the slice data to the decoding stage.

Data parser tracks the length of the slice, which is provided by the MEDIA_OBJECT command. Data parser uses the slice length and the starting offsets to calculate the end of slice. When the end of slice is reached, data parser indicates end of slice to symbol decoder and does not pass on any more data that comes from the command stream until a new slice begins.



Symbol Decoder

Symbol decoder performs variable length decoding of the slice bitstream according to the MPEG-2 standard. The decoder analyzes symbols in the bitstream and separates them for further processing. For example, motion vector differentials are sent to motion vector predictor but DCT coefficients are sent directly to IS stage.

Motion Vector Predictor

Motion Vector (MV) Predictor calculates the motion vectors based on the motion vector differentials received from symbol decoder and the motion vector prediction values maintained within MV Predictor, updates the motion vector prediction values accordingly and performs additional arithmetic for dual prime motion vectors to convert them to uni/bi-directional motion vectors. The output motion vectors are relative to the current macroblock position.

1.4.2.3.2 Inverse Scan

IS unit process the non-zero DCT coefficients with their (x, y) location within an 8x8 block received from VLD on a block by block basis. For each new block of data, IS initializes the 8x8 block storage to zero. For each non-zero coefficient received from VLD, IS first sign-extend it to a 16-bit signed value and then place it in the block storage at the location identified by its (x, y) address. When the end of block signal is received from VLD, IS writes the assembled DCT data block into URB.

Only the coded blocks are assembled in the URB, and they are assembled back to back. As the thread payload size is variable, VFE calculates the final message size according to the coded block pattern field before sending the payload size to TS.

1.4.2.3.3 Output Formatting

Additional functionality after inverse scan formats the data that is sent as thread payload to the kernel. Some of this functionality, such as expansion of skip macroblocks and determination of second P field, is done by hardware to make the kernel more efficient.

Skip Macroblocks

VFE processes skip macroblocks by separating them into individual macroblocks and forming one thread for one macroblock. For each skip macroblock, hardware sets its coded block pattern to 0, indicating that no error data is present. All contiguous skip macroblocks have the same relative motion vector. Hardware also handles the difference of skipped macroblocks in a P picture or a B picture as defined by MPEG-2 specification. According to MPEG-2 specification, skip macroblocks cannot extend beyond the end of the current line.

Second Field

According to MPEG-2 specification, field prediction for a P field picture uses the most recently decoded two fields as reference, namely, the most recently decoded reference top field and the most recently decoded reference bottom field. As shown in **Error! Reference source not found.**, when the current P field is the first field of a frame, both of its reference fields come from the same frame. When the current P field is the second field of a frame, one of its reference fields comes from the same frame. This is illustrated in **Error! Reference source not found.**

Detecting second field is important if reference frame selection is required. This is no longer true for GEN4 as each reference field is specified by unique binding table index. Each binding table index contains the pointer to the surface state, which contains not only the field indication but also the base address of the frame buffer. Therefore, it is up to the kernel developer to determine whether to use the Second Field information provided by the hardware.

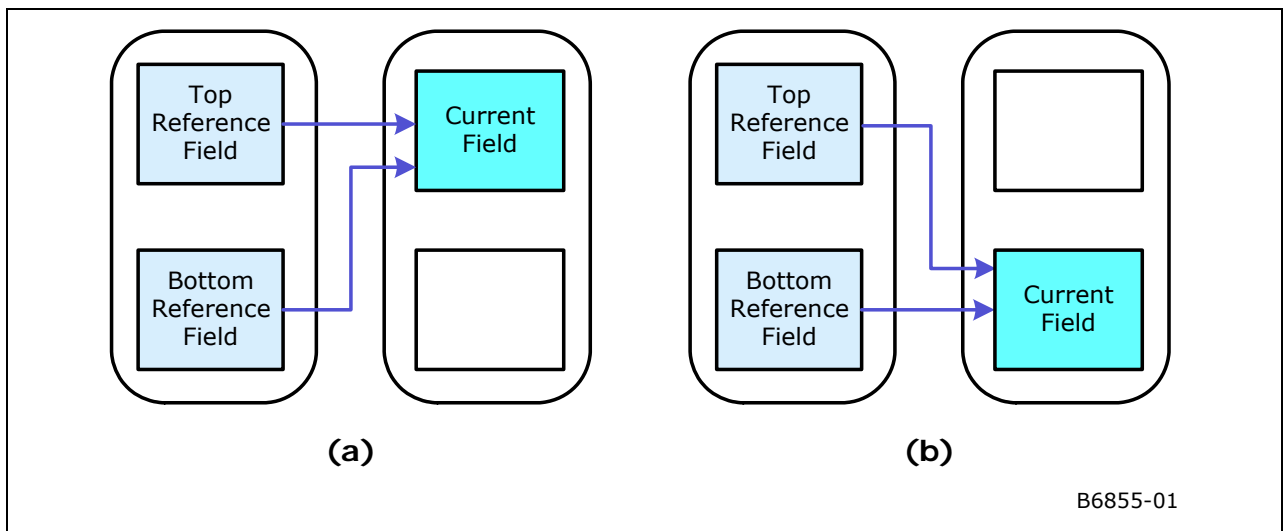
VFE sets the second field indicator under the following conditions:

- Picture coding type is P picture
- Destination format is field
- Motion type may be field, 16x8 or dual prime
- Either:
 - Top field is first, and
 - Current field is field 1, and

OR

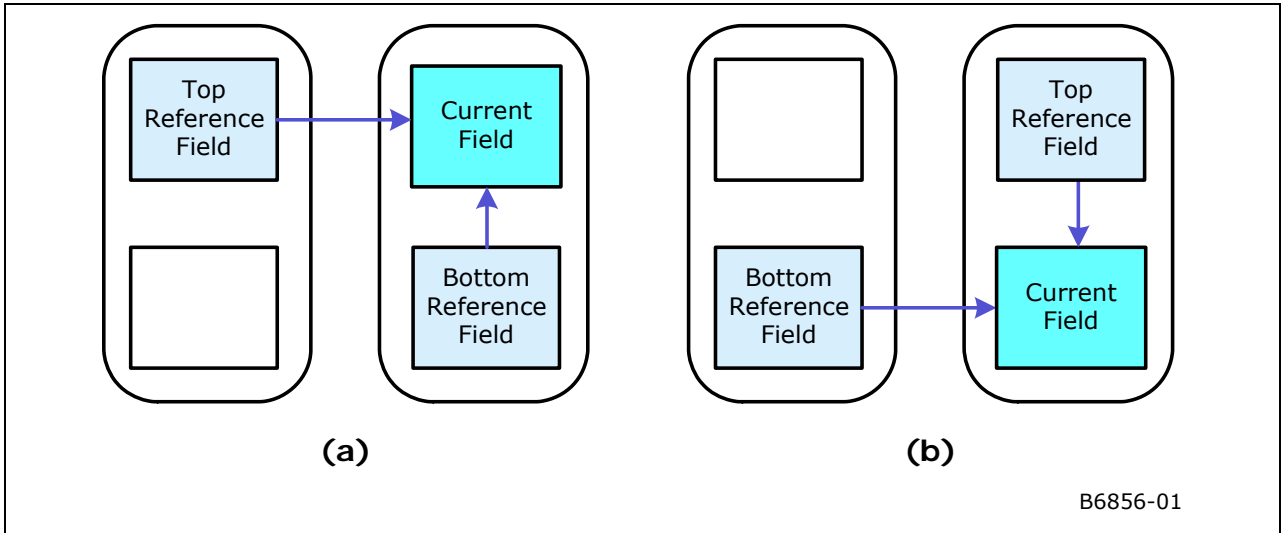
- Top field is not first, and
- Current field is field 0, and

Prediction for a P field picture that is a first field, which is (a) a top field, or (b) a bottom field.





Prediction for a P field picture that is a second field, which is (a) a top field, or (b) a bottom field





1.4.2.3.4 Handling Motion Vectors

Error! Reference source not found. provides a summary of different motion types and associated properties. For Frame_Motion_Type, there are three types of Prediction_Type: frame-based prediction, field-based prediction and dual-prime prediction. For Field_Motion_Type, there are three types of Prediction_Type: field-based prediction, dual-prime prediction and 16x8 prediction.

The table below details the motion compensation operations for various frame motion types and **Error! Reference source not found.** depicts the motion compensation operations for various field motion types.

Summary of Motion Types

| *_Motion_Type | Prediction_Type | Vector [r][s] | Possible MV Combinations in Bitstream | Uses Motion Vertical Field Select |
|---------------|-----------------|--|---------------------------------------|-----------------------------------|
| Frame | Frame-based | [0][0] – 0 [0][1] – 1 [1][0] – 2 [1][1] – 3 | None,0,1,0+1 | No |
| Frame | Field-based | [0][0] – 0 [0][1] – 1 [1][0] – 2 [1][1] – 3 | 0+2,1+3, 0+1+2+3 | Yes |
| Frame | Dual-Prime | [0][0] – 0 [0][1] – 1 [1][0] – 2 [1][1] – 3 [2][0] – 4 [3][0] – 6 | 0+4+6 (See Frame-Dual Prime table) | No |
| Field | Field-based | [0][0] – 0 [0][1] – 1 [1][0] – 2 [1][1] – 3 | None, 0, 1, 0+1 | Yes |
| Field | Dual-Prime | [0][0] – 0 [0][1] – 1 [1][0] – 2 [1][1] – 3 [2][0] – 4 | 0+4 (See Field-Dual Prime table) | Yes |
| Field | 16x8 | [0][0] – 0 [0][1] – 1 [1][0] – 2 [1][1] – 3 | 0+2, 1+3,0+1+2+3 | Yes |

*Vectors 4 and 6 are the derived motion vectors (DMVs) for dual-prime prediction that are calculated by PR and placed in the thread payload in the specified MVector position.



Motion Comp Operation for Pictures with Frame Motion Type

| frame_motion_type | forward | backward | intra | Motion vector (v[r][s][t]) | Command | HW Mvector (MV[r][s]) | Prediction Map | Prediction formed for | MVFS |
|-------------------|----------|----------|-------|----------------------------|---------|-----------------------|----------------|---|------------|
| Frame-based‡ | - | - | 1 | v[0][0][1:0] | 0 | - | - | None (motion vector is for concealment) | - |
| Frame-based | 1 | 1 | 0 | v[0][0][1:0] | 2 | MV[0][0] | Fwd | frame, forward | - |
| | | | | v[0][1][1:0] | | MV[0][1] | Back | frame, backward | - |
| Frame-based | 1 | 0 | 0 | v[0][0][1:0] | 2 | MV[0][0] | Fwd | frame, forward | - |
| Frame-based | 0 | 1 | 0 | v[0][1][1:0] | 2 | MV[0][1] | Back | frame, backward | - |
| Frame-based‡ | 0 (1) | 0 | 0 | v[0][0][1:0]*§ | 2 | MV[0][1] | Fwd | frame, forward | - |
| Field-based | 1 | 1 | 0 | v[0][0][1:0] | 4 | MV[0][0] | Fwd | top field, forward | [0][0] |
| | | | | v[1][0][1:0] | | MV[1][0] | Fwd | bottom field, forward | [1][0] |
| | | | | v[0][1][1:0] | | MV[0][1] | Back | top field, backward | [0][1] |
| | | | | v[1][1][1:0] | | MV[1][1] | Back | bottom field, backward | [1][1] |
| Field-based | 1 | 0 | 0 | v[0][0][1:0] | 4 | MV[0][0] | Fwd | top field, forward | [0][0] |
| | | | | v[1][0][1:0] | | MV[1][0] | Fwd | bottom field, forward | [1][0] |
| Field-based | 0 | 1 | 0 | v[0][1][1:0] | 4 | MV[0][1] | Back | top field, backward | [0][1] |
| | | | | v[1][1][1:0] | | MV[1][1] | Back | bottom field, backward | [1][1] |
| Dual prime | 1 | 0 (1) | 0 | v[0][0][1:0] | 4 | MV[0][0] | Fwd | top field, from same parity, forward | [0][0] = 0 |
| | | | | v[0][0][1:0] | | MV[1][0] | Fwd | bottom field, from same parity, forward | [1][0] = 1 |
| | | | | v[2][0][1:0]*† | | MV[0][1] | Fwd | top field, from opposite parity, forward | [0][1] = 1 |
| | | | | v[3][0][1:0]*† | | MV[1][1] | Fwd | bottom field, from opposite parity, forward | [1][1] = 0 |

NOTE -
 ? the motion vector is only present if concealment_motion_vectors is one
 ‡ frame_motion_type is not present in the bitstream but is assumed to be Frame-based
 * These motion vectors are not present in the bitstream
 † These motion vectors are derived from vector v[0][0][1:0] as described in 7.6.3.6
 § The motion vector is taken to be (0; 0) as explained in 7.6.3.5



Motion Comp Operation with Field Motion Type

| field_motion_type | forward | backward | intra | Motion vector | Command | Prediction Map (SecondPField BottomField) | | | | Prediction formed for | | |
|------------------------|---------|----------|-------|--------------------------|----------|---|----------|----------|---------------------------|---------------------------|--|--|
| | | | | | | HW MVector | 00 | 01 | 10 | | | 11 |
| Field-based‡ | - | - | 1 | v'[0][0][1:0]? | None | N/A | - | - | - | - | None (motion vector is for concealment) | |
| Field-based | 1 | 1 | 0 | v'[0][0][1:0] | 2 | MV[0][0] | Fwd | Fwd | x | x | whole field, forward | B-Pict only |
| | | | | v'[0][1][1:0] | MV[0][1] | Back | Back | x | x | whole field, backward | | |
| Field-based, P picture | 1 | 0 | 0 | v'[0][0][1:0] | 2 | MV[0][0] | Fwd (M0) | Fwd (M0) | Fwd (M0) if MVFS[0][0]=0 | Dst (M1) if MVFS[0][0]=0 | whole field, forward | |
| | | | | | | | | | Dst (M1) if MVFS[0][0]=1 | Fwd (M0) if MVFS[0][0]=1 | | |
| Field-based, B picture | 1 | 0 | 0 | v'[0][0][1:0] | 2 | MV[0][0] | Fwd | Fwd | x | x | whole field, forward | |
| Field-based | 0 | 1 | 0 | v'[0][1][1:0] | 2 | MV[0][1] | Back | Back | x | x | whole field, backward | B-Pict only |
| Field-based‡ | 0 (1) % | 0 | 0 | v'[0][0][1:0]*§ | 2 | MV[0][0] | Fwd | Fwd | Fwd w/ MVFS[0][0]=0 | Fwd w/ MVFS[0][0]=1 | whole field, forward | P-Pict only, Same parity. |
| 16x8 MC | 1 | 1 | 0 | v'[0][0][1:0] | 4 | MV[0][0] | Fwd | Fwd | x | x | upper 16x8 field, forward | B-Pict only |
| | | | | v'[1][0][1:0] | | MV[1][0] | Fwd | Fwd | x | x | lower 16x8 field, forward | |
| | | | | v'[0][1][1:0] | | MV[0][1] | Back | Back | x | x | upper 16x8 field, backward | |
| | | | | v'[1][1][1:0] | | MV[1][1] | Back | Back | x | x | lower 16x8 field, backward | |
| 16x8 MC | 1 | 0 | 0 | v'[0][0][1:0] | 4 | MV[0][0] | Fwd | Fwd | Fwd (M0) if MVFS[0][0]=0 | Dst (M1) if MVFS[0][0]=0 | upper 16x8 field, forward | |
| | | | | Dst (M1) if MVFS[0][0]=1 | | Fwd (M0) if MVFS[0][0]=1 | | | | | | |
| | | | | v'[1][0][1:0] | | MV[1][0] | Fwd (M0) | Fwd (M0) | Fwd (M0) if MVFS[1][0]=0 | Dst (M1) if MVFS[1][0]=0 | lower 16x8 field, forward | |
| | | | | Dst (M1) if MVFS[1][0]=1 | | Fwd (M0) if MVFS[1][0]=1 | | | | | | |
| 16x8 MC | 0 | 1 | 0 | v'[0][1][1:0] | 4 | MV[0][1] | Back | Back | x | x | upper 16x8 field, backward | B-Pict only |
| | | | | v'[1][1][1:0] | | MV[1][1] | Back | Back | x | x | lower 16x8 field, backward | |
| Dual prime | 1 | 0 (1) % | 0 | v'[0][0][1:0] | 2 | MV[0][0] | Fwd (M0) | Fwd (M0) | Fwd (M0) w/ MVFS[0][0]=0 | Fwd (M0) w/ MVFS[0][0]=1 | whole field, from same parity, forward | P-Pict only (SW forces MVFS[0][0], MVFS[0][1]) |
| | | | | v'[2][0][1:0]*† | | MV[0][1] | Fwd (M1) | Fwd (M1) | Dest (M1) w/ MVFS[0][1]=1 | Dest (M1) w/ MVFS[0][1]=0 | whole field, from opposite parity, forward | |

Notes: Motion vectors are listed in the order they appear in the bitstream.

? — The motion vector is only present if concealment_motion_vectors is one.

‡ — field_motion_type is not present in the bitstream but is assumed to be Field-based.

* — These motion vectors are not present in the bitstream.

† — These motion vectors are derived from vector v'[0][0][1:0] as described in 7.6.3.6.

§ — The motion vector is taken to be (0; 0) as explained in 7.6.3.5.

% — Software converts the motion type. For Dual-prime case, software converts it to a bidirectional prediction.



Remarks:

- Forward, Backward MV pairs always used for combined prediction (Bi-dir or Dual-Prime)
- MVs [0][0] and [0][1] can be to Fwd reference, Backward reference, or Destinsation buffer (for 2nd field case)
- MV [1][0] can be to a Fwd or Dest
- MV [1][1] is always to a Back
- (M0), (M1) stands for MIP_INFO setting for the first reference picture and the second reference picture. It is particularly important to set correct M1 for P-pictures to deal with SecondField and DualPrime cases.
- Software converts the dual-prime case to a field-based bidirectional prediction with 2 MVs.

1.4.2.3.5 Dual Prime Handling

Dual prime prediction is only valid for a P-picture. In dual prime mode, each field will have two predictions similar to the forward and backward predictions in a B-picture, as the final prediction value for the field is the average of the two. One of the motion vectors is provided by the bitstream and the other one is derived. Motion Vector Predictor unit is responsible for converting all dual prime predictions to a forward and backward field prediction according the **Error! Reference source not found.** for P frame picture and **Error! Reference source not found.** for P field picture.

Converting Frame-Dual Prime Motion to 4MV

| Prediction formed for: Field / Parity | MPEG-2 MV[r][s] | Thread Payload MV[r][s] | Motion Vertical Field Select (MVFS) | |
|--|--------------------|----------------------------|-------------------------------------|-----------|
| Top / Same | [0][0] | [0][0] | Bit 0 = 0 | MVFS = 6h |
| Bottom / Same | [0][0] | [1][0] | Bit 2 = 1 | |
| Not Used | [0][1] | - | | |
| Not Used | [1][0] | - | | |
| Not Used | [1][1] | - | | |
| Top / Opposite | [2][0] | [0][1] | Bit 1 = 1 | |
| Bottom / Opposite | [3][0] | [1][1] | Bit 3 = 0 | |

Converting Field-Dual Prime Motion to 2MV

| Prediction formed for: Field / Parity | MPEG-2 MV[r][s] | Thread Payload MV[r][s] | MotionVerticalFieldSelect | |
|--|--------------------|----------------------------|---------------------------|--------------|
| | | | Top Field | Bottom Field |
| Whole field / Same | [0][0] | [0][0] | Bit 0 = 0 | Bit 0 = 1 |
| Whole field / Opposite | [2][0] | [0][1] | Bit 1 = 1 | Bit 1 = 0 |



1.4.2.3.6 Interface Descriptor Selection

In VLD mode, the Interface Descriptor Offset field in the MEDIA_OBJECT command is ignored by hardware. Instead, the interface descriptor offset is computed by hardware based on the decoded macroblock parameters and a remapping table.

First a macroblock index is computed based on parameters such as picture structure, motion type, prediction type, DCT type, intra-coding type and motion vector present information. **Error! Reference source not found.** provides the macroblock index table for a frame-picture destination buffer (with Picture Structure = 11). **Error! Reference source not found.** shows macroblock indices for a field-picture destination buffer (with Picture Structure = 01 or 10). As Picture Structure is a state variable that will not be changed until a pipeline flush, the macroblock indices can be computed separately for different Picture Structure.

After the macroblock index is computed, it is used as the index into the Interface Descriptor Remap Table to derive the final interface descriptor offset value. The Interface Descriptor Remap Table is provided as part of the VLD state.

The interface descriptor offset value multiplied by the interface descriptor size is then added to the interface descriptor base pointer to generate the interface descriptor pointer for the post-VLD thread.

The last three columns in **Error! Reference source not found.** and **Error! Reference source not found.** indicate whether a macroblock index is applicable for a given Picture Coding type (I, P or B). A ‘Y’ (or a ‘N’) means the macroblock index on the row is valid (or invalid) for the Picture Type shown on the column. Taking a frame picture destination for example, only macroblock indices 0 and 8 are valid for an I-picture; indices 0-3 and 8-11 are valid for a P-picture; and for a B-picture, only indices 3 and 11 are not valid.

Developers can use the remap table for kernel development to fine-tune system performance and reduce software complexity. For example, if the destination is a frame picture, the kernel for a macroblock with dual-prime motion in a P-picture (macroblock index = 3) may be identical to that for a macroblock with bidirection field motion in a B picture (macroblock index = 7). A common set of interface descriptors can be configured once for frame picture destinations, and reused without change when the destination is of I-, P- and B- picture coding type.

In another case, if it is determined that kernel software will be responsible of handling DCT types for a frame picture destination, then macroblock index i and $i+8$, for $i = 0$ to 7 , can be mapped to the same interface descriptor.

Macroblock indices for frame picture destination

| Macroblock Index | Interface Descriptor Kernel Function (Frame Picture Destination) | I | P | B |
|------------------|--|---|---|---|
| 0 | I macroblock | Y | Y | Y |
| 1 | Forward frame motion | N | Y | Y |
| 2 | Forward field motion | N | Y | Y |
| 3 | P picture, dual-prime motion | N | Y | N |
| 4 | Backward frame motion | N | N | Y |
| 5 | Backward field motion | N | N | Y |
| 6 | Bidirectional frame motion | N | N | Y |
| 7 | Bidirectional field motion | N | N | Y |
| 8 | I macroblock w/ field DCT | Y | Y | Y |
| 9 | Forward frame motion w/ field DCT | N | Y | Y |



| Macroblock Index | Interface Descriptor Kernel Function (Frame Picture Destination) | I | P | B |
|------------------|--|---|---|---|
| 10 | Forward field motion w/ field DCT | N | Y | Y |
| 11 | P picture, dual-prime motion w/ field DCT | N | Y | N |
| 12 | Backward frame motion w/ field DCT | N | N | Y |
| 13 | Backward field motion w/ field DCT | N | N | Y |
| 14 | Bidirectional frame motion w/ field DCT | N | N | Y |
| 15 | Bidirectional field motion w/ field DCT | N | N | Y |

Macroblock indices for field picture destination


| Macroblock Index | Interface Descriptor Kernel Function (Field Picture Destination) | I | P | B |
|------------------|--|---|---|---|
| 0 | I macroblock | Y | Y | Y |
| 1 | Forward field motion | N | Y | Y |
| 2 | Forward 16x8 motion | N | Y | Y |
| 3 | P picture, dual-prime motion | N | Y | N |
| 4 | Backward field motion | N | N | Y |
| 5 | Backward 16x8 motion | N | N | Y |
| 6 | Bidirectional field motion | N | N | Y |
| 7 | Bidirectional 16x8 motion | N | N | Y |

1.4.2.4 VC1-IT Mode [DevCTG, DevILK]

In VC1-IT mode, Inverse Transform is performed by VFE hardware and Motion Compensation is done by GEN4 kernels. Overlap Smoothing is performed jointly by VFE hardware and GEN4 kernels.

1.4.2.4.1 Overlap Smoothing

One unique feature required by VC1 standard is the Overlap Smoothing operation. Overlap Smoothing involves filtering of the edges of two neighbor Intra blocks after Interface Transform. A hybrid solution is employed in GEN4 architecture with vertical edge filtering done by VFE hardware and horizontal edge filtering done by kernels. This is illustrated in **Error! Reference source not found.** where 'A' and 'B' are two (8x8) blocks. The shaded areas around Intra block 'B' are filtered if the corresponding neighbor block is also an Intra block. Information regarding whether an edge requires filter is carried in the VC1-IT in-line data.



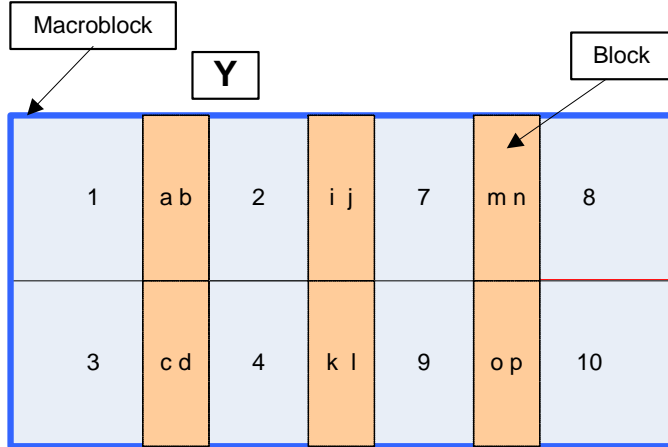
Overlap Smoothing along the edges of neighboring Intra blocks

As each MEDIA_OBJECT_EX command corresponds to one macroblock, vertical edge filtering within a macroblock can be performed in stream by VFE. For edges cross macroblocks, VFE stores a 2x8 strip for each block along the macroblock right-hand-side edges and uses the stored strips for blocks from the next macroblocks.


Error! Reference source not found. shows an example of processing two adjacent macroblocks. The blocks are delivered to VFE in order shown below from 1 to 12. The shaded strips, namely 'a' to 'p', are 2x8 block boundaries that need to be filtered. The edges within a macroblock boundary are filtered when Inverse Transform is processed for the current macroblock. The edges of a macroblock are stored in VFE and modified when the next macroblock is transformed and then written to the URB. Note that VFE stores four right-most block edge strips for the current macroblock (*i,k,e,g*).

In summary, VFE hardware has a delayed operation of Overlap Smoothing such that at the end, each block is still output in 8x8 quantity with its left and right edges filtered. Storing such edge strips means that VFE maintain an internal state, which is not saved or restored during context switch. This has the following implications


- Software must deliver the macroblocks in row-major order (the natural macroblock decoding order) whenever Overlap Smoothing is turned on.
- Software must avoid context switch between two Intra macroblocks.



Cb




Cr



Overlap Smoothing vertical filter operation by VFE


The Overlap Smoothing for horizontal edges is performed by kernels in the EU. Indicators of the edges to be smoothing are sent as control in the in-line data. As higher precision (16-bit) are required to store the edge pixels comparing to the 8-bit precision for the final destination, a separate surface must be used.

As shown by the example in **Error! Reference source not found.**, when processing a macroblock that contains the block with pixels in region 'A' and region 'x', the kernel may write the final pixel data (in 8-bit precision) in region A to the destination surface, while keeping a copy of the pixel data in region 'x' (in 16-bit precision) to a difference buffer, as pixels in region 'x' may or may not require Overlap Smoothing depending on future macroblocks containing pixels in regions 'y' and 'B'. Later on, Macroblock containing region B is processed. If Overlap Smoothing is enabled across the horizontal edges between A and B, the kernel (processing B) reads data from region x, performs overlap smoothing on x and y blocks and writes macroblock B (including region 'y') to the destination buffer and updates the two rows for macroblock A as well. The LastRowFlag in the VC1-IT inline data can be used by kernel to know if the current macroblock is in the last row and therefore should write the last two rows of pixels to the destination surface.



Overlap Smoothing horizontal filter operation by kernels

In addition, kernels must also take care of the thread-to-thread dependencies between vertically adjacent macroblocks. As shown in XXX, if macroblock at location (X,Y') requires Overlap Smoothing filtering on its upper edge, it can not proceed Overlap Smoothing operation until the thread operating on macroblock at (X, Y) has terminated with output data reached coherent space.



Dispatch X,Y' only after
X,Y is complete from EU

Thread-thread Dependency in Overlap Smoothing horizontal filter operation by kernels

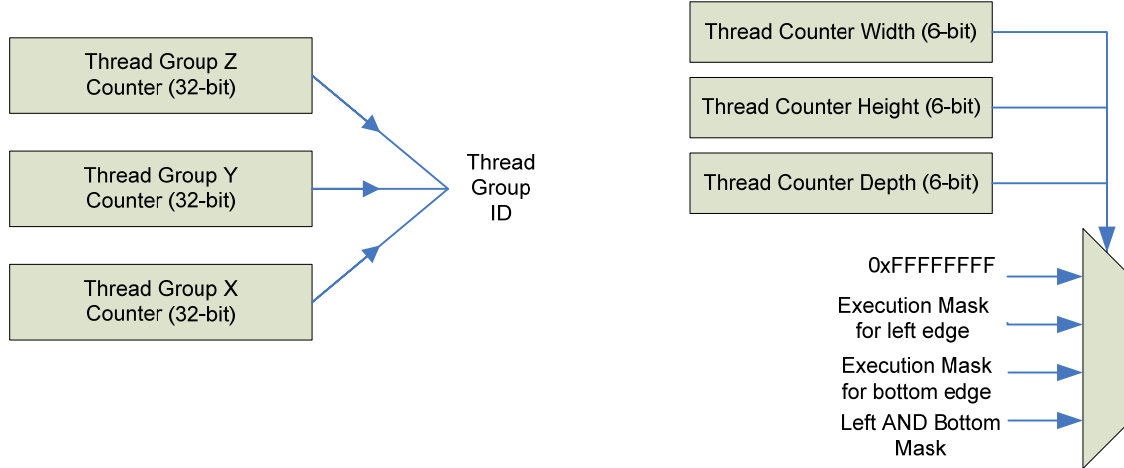
1.4.2.5 GPGPU Mode

The general purpose (GPGPU) mode allows the Gen7 architecture to be used by general purpose parallel APIs such as OpenCL and the DX11 ComputeShader. This is similar to the Generic mode with additional support for automatic generation of threads, Shared Local Memory and Barriers.

1.4.2.5.1 Automatic Thread Generation

A single GPGPU job may require thousands or even millions of GPU_OBJECT commands. Rather than create them separately, it would be better to generate them algorithmically. To do this a GPGPU_WALKER command is created.

Rather than modifying the Media Walker, a simple Thread Group Walker is created instead:



The X/Y/Z counters for the thread group will have an initial and maximum value. The thread group id sent with each dispatch consists of these 3 numbers. These counters are 32-bits since the spec does not give a limit to the size of the thread id.

The 3 thread counters count the number of dispatches in a single thread group – up to 32 dispatches for SIMD32 or 64 dispatches for SIMD16/8. There are 3 of them in order to select the execution masks correctly – see section 1.4.2.5.3 on execution masks. Each one is 6-bits in order to allow full flexibility of any dimension going to 64 while the rest do not increment.

A thread is generated each time one of the thread counters increment. When all the counters are reach thier maximum value the thread group is done and the thread group counter can increment and start a new thread group. When the thread group X counter reaches it's maximum it is reset to it's initial value and the Y counter is incremented.

The compiler determines how many SIMD channels are needed per thread group, and then decides how these will be split among EU threads. The number of threads is programmed in the thread counter, and the SIMD mode (SIMD8/SIMD16/SIMD32) is specified in the GPGPU_WALKER command.

1.4.2.5.2 Thread Payload

The payload to each thread dispatched is:

- 1) A thread group id which identifies the group the set of threads belong to. This is in the form of a set of 3, 32-bit X/Y/Z values.
- 2) The set of X/Y/Z that form the thread id for each channel. If Z is not used then only X/Y are needed.
- 3) The execution mask which indicates which channels are active.

Thread ids form a 2D or 3D surface which has to be mapped into SIMD32, SIMD16 or SIMD8 dispatches. Rather than have the hardware force a particular mapping of thread ids to channels, the



mapping will be supplied by the compiler. The VFE will receive a simple count of the number of threads per thread group which will be used to count the number of dispatches. The thread ids for all threads in a thread group are put in a constant buffer with the MEDIA_CURBE_LOAD command. A single set of thread ids can be used repeatedly for all thread groups, since the thread ids are the same for each thread group id output by the GPGPU_WALKER.

The data required is up to the compiler, but here is an example set of payloads for a 2 Z x 2Y x 12 X and a SIMD16 dispatch. This thread group requires 3 dispatches:

| | |
|-------------------------------------|---------------------------|
| 3 2 1 0 11 10 9 8 7 6 5 4 3 2 1 0 | Thead id X for dispatch 0 |
| 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 | Thead id Y for dispatch 0 |
| 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 | Thead id Z for dispatch 0 |
| 7 6 5 4 3 2 1 0 11 10 9 8 7 6 5 4 | Thead id X for dispatch 1 |
| 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 1 | Thead id Y for dispatch 1 |
| 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 | Thead id Z for dispatch 1 |
| 11 10 9 8 7 6 5 4 3 2 1 0 11 10 9 8 | Thead id X for dispatch 2 |
| 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 | Thead id Y for dispatch 2 |
| 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 | Thead id Z for dispatch 2 |

In this case the thread counter width would be programmed with a maximum value of 3 (since all the execution masks are all F, it doesn't matter how the thread counters are programmed as long as they count to 3 before finishing the thread group).

The first dispatch would tell the TS (who would tell the TD) that the payload starts at the beginning of the constant buffer and has a length of 3. The next dispatch would have a payload starting at constant_buffer_start + 3. The final dispatch payload starts at constant_buffer_start + 6. If there are more thread groups in the command they would get exactly the same payload – the only difference is the thread group id (as well as a different barrier and shared local memory space).

1.4.2.5.3 Execution Masks

The number of channels required by the GPGPU job may not evenly fit into the number of SIMD channels. That can leave some channels idle. The execution mask is used to tell the hardware which channels are to be used.

A thread group is modeled as a 3D solid with each channel acting as one X/Y/Z point in the solid. This can take the form of a line with 1024 channels with X from 0 to 1023 and constant Y/Z, a square with X=0 to 32 and Y=0 to 32, or a cube with X=0 to 9, Y=0 to 9, Z=0 to 9. Software needs to determine how these shapes are mapped onto the 32 SIMD32 channels per dispatch (or 16 SIM16, etc). The mapping per thread is assumed to be a 2D square of channels such as 8x4, 16x2, 32x1. Below is a diagram of a 22x6 thread group that is mapped onto a set of 8x4 SIMD32 channels:




| | | | |
|-----|-----|-----|--|
| 8x4 | 8x4 | 8x4 | |
| 8x4 | 8x4 | 8x4 | |
| 8x4 | 8x4 | 8x4 | |
| | | | |

Note that the dispatches to the top and left have execution masks of all-F, while all the right edge dispatches have the same execution mask; likewise all the bottom edge dispatches have the same execution mask. The bottom right is the logical-AND of the right and bottom edge dispatches.

A 32-bit right and bottom mask is sent with the GPGPU_WALKER command, and the thread width, height and depth counters are used to determine when they are used (width, height and depth are used instead of X/Y/Z, since it is not required that width = X – width and height are the two variables that are changing in a single SIMD dispatch even if they are Y and Z).

For each dispatch the width counter is incremented until it reaches the maximum – the dispatch with width=max will use the right execution mask. The height counter is then incremented and process repeated. If at any time the height counter = max then the execution mask is the bottom execution mask. When the height and width counters are both max then the dispatch will be the AND of the right and bottom and the depth counter will increment.

The same 2Z x 2Y x 12X thread group described above dispatched as SIMD32 with each dispatch delivering a 16X x 2Y shape would require 2 dispatches with empty bits in the right execution mask and all F in the bottom.



The width and height counter would have a maximum of 1, and the depth counter would have a maximum of 2. The two dispatches would use the AND of the two masks, but since the bottom mask is F it would be the same as just the right mask.



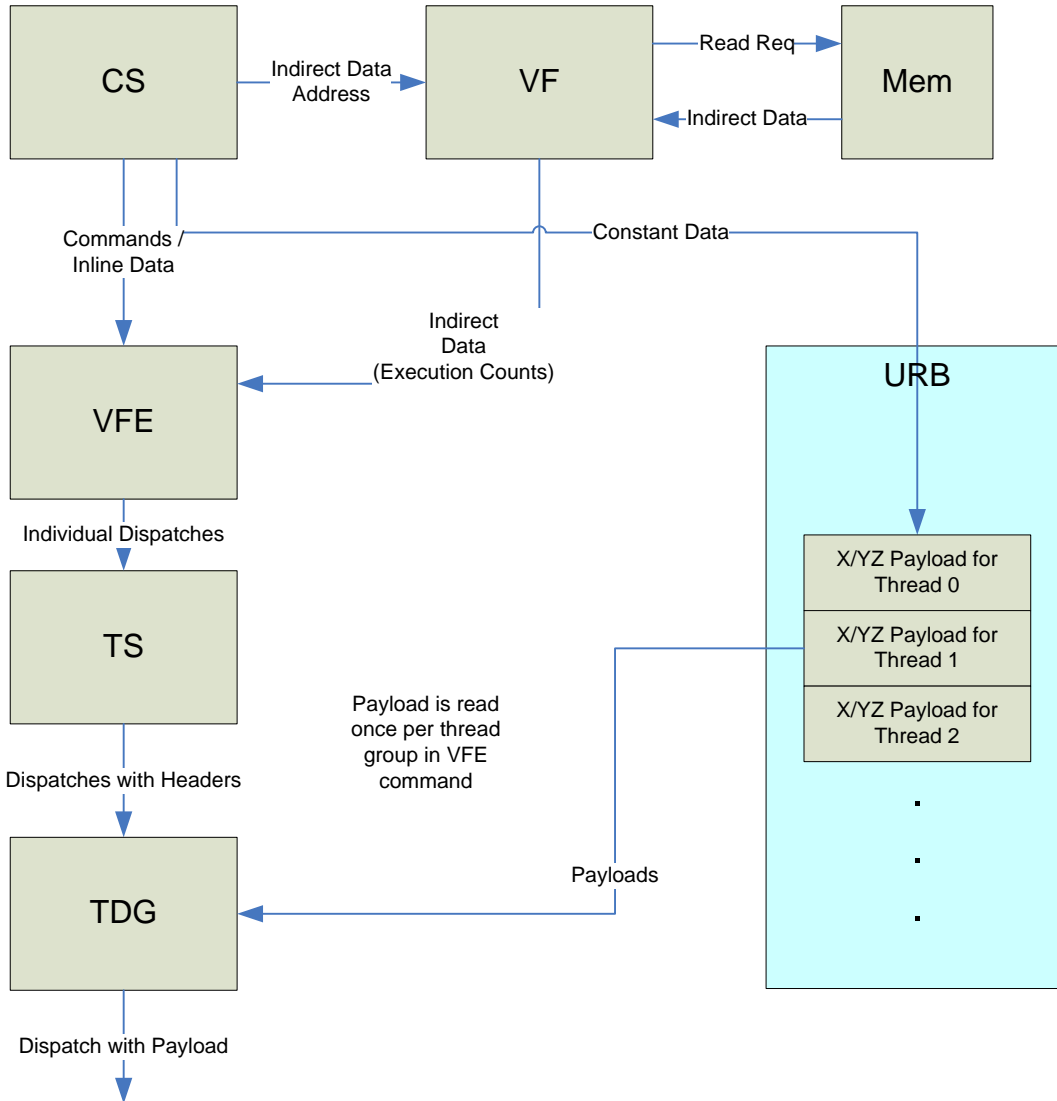
The execution masks can also be used when the software wants to pack the channels rather than lay them out in a regular pattern:

| | |
|-------------------------------------|---------------------------|
| 3 2 1 0 11 10 9 8 7 6 5 4 3 2 1 0 | Thead id X for dispatch 0 |
| 7 6 5 4 3 2 1 0 11 10 9 8 7 6 5 4 | |
| 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 | Thead id Y for dispatch 0 |
| 0 0 0 0 0 0 0 0 1 1 1 1 1 1 1 | |
| 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 | Thead id Z for dispatch 0 |
| 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 | |
| 11 10 9 8 7 6 5 4 3 2 1 0 11 10 9 8 | Thead id X for dispatch 1 |
| | |
| 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 | Thead id Y for dispatch 1 |
| | |
| 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 | Thead id Z for dispatch 1 |
| | |

In this case the width counter can have a maximum of 2, and the height and depth counters with a maximum of 1. The first dispatch will use the bottom mask only (all-F) and the second will use the right AND bottom mask to remove the channels that are not used.

1.4.2.5.4 Payload Storage

The MEDIA_CURBE_LOAD constant data is stored in the URB by CS and read out by TDL when the dispatch occurs. The inline payload with the execution counts is sent to VFE from CS, indirect payload is read by VF and sent to VFE. The execution counts are stored internal to VFE.



Only 32 threads are allowed for SIMD32 to match the 1024 thread limit, requiring 32 execution count bytes, or 8 DW payload. SIMD16 and SIMD8 allow the full 64 thread per half-slice, and so require as much as 16 Dwords.

The X/Y/Z payload size per dispatch is specified in the command, but a maximum size is 3 16-bit numbers per 1024 SIMD channels, or 6 kbytes.

1.4.2.5.5 URB Management

The VFE manages the URB in GPGPU and generic/media modes. The first 32 URB entries are reserved for the interface descriptor, and CURBE data is placed after the IDs. URB handles are needed for indirect data and parent/child communication; when the VFE starts up it will create up to 64 handles by partitioning the remaining URB space into evenly spaced addresses and saving



the resulting handles in a FIFO. The handles can then be treated just like ones created by the URBM – send to TD on dispatch and recovered on the handle return bus.

MEDIA_VFE_STATE specifies the amount of CURBE space, the URB handle size and the number of URB handles. The driver must ensure that $((\text{URB_handle_size} * \text{URB_num_handle}) - \text{CURBE} - 32) \leq \text{URB_allocation_in_L3}$.

1.4.3 Scoreboard Control [DevILK+]

A hardware mechanism is provided to control the dispatch of root threads. Without using this hardware mechanism, only the dispatch of a SRT is managed by a parent root thread using the SRT message to TS.

There is a scoreboard hardware in TS unit. The scoreboard is addressed by the 18-bit (X, Y) scoreboard field in VFE DWord, where (X, Y) is typically used as the Cartesian coordinate of the working unit in a 2D frame but may be interpolated in other ways. When a root thread is dispatched, the entry at (X, Y) is marked. When the root thread is terminated, the corresponding bit in the scoreboard is cleared.

Each root thread may have up to eight dependencies. The dependency relation is described by the state value of Scoreboard Controls in terms of related distance of (deltaX, deltaY). There is a global scoreboard enabling in the state as well as the-per thread enabling for each dependency.

TS stalls the dispatch of a root thread if any scoreboard entry, which is denoted by (Scoreboard X + deltaX, Scoreboard Y + deltaY), matching with any enabled dependencies is marked as in-flight. The thread is dispatched only after all dependencies are cleared.

For a root thread, TS only stalls the dispatch of the thread only if the dependent scoreboard entries of the thread marked. It does not automatically stalls the dispatch for destination collision if (deltaX = 0, deltaY=0) is not set in the scoreboard state. This kind of scoreboard destination collision is due to the scoreboard wrap-around (or aliasing), which must be avoided. With 9-bit per X, Y field, the hardware scoreboard can support a frame that is subdivided up to 512x512 threads without a scoreboard aliasing.

In addition to the above ‘stalling scoreboard’, Media Pipe may also support a non-stalling scoreboard. With the non-stalling, a thread is dispatched with the dependent threads marked. The thread dependency affects the issuing of a SENDC instruction. See ISA chapter for details.

Scoreboard Support in Device Hardware

| Device | Stalling scoreboard | Non-Stalling scoreboard |
|----------|---------------------|-------------------------|
| [DevILK] | Yes | No |

Restrictions:

- *The hardware scoreboard only handles root threads, but not child threads. This limitation may be revisited when future application requirement changes.*
- *The usage of hardware scoreboard and SRT are maturely exclusive. In other words, when hardware scoreboard is used, SRT should not be issued.*

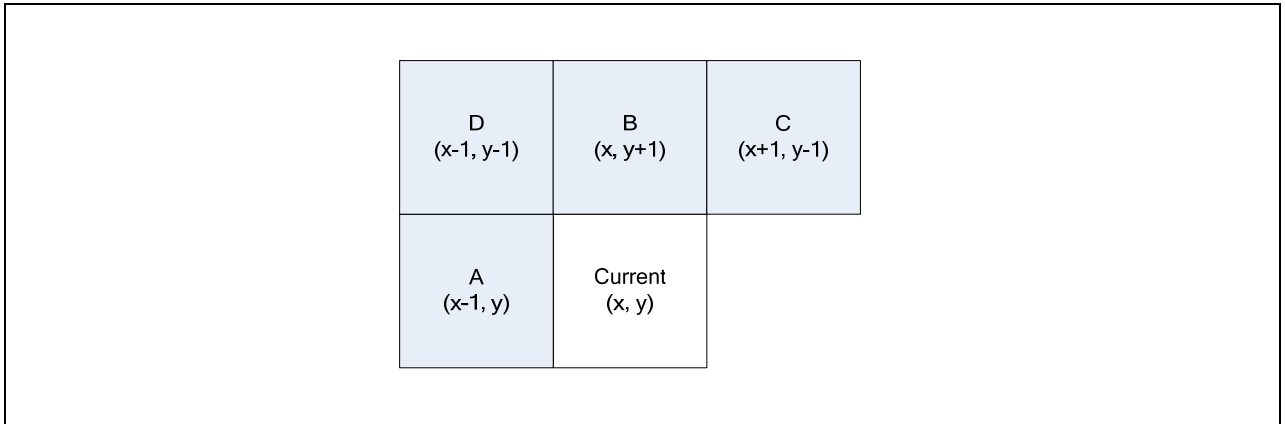


1.4.3.1 AVC-Style Dependency Example

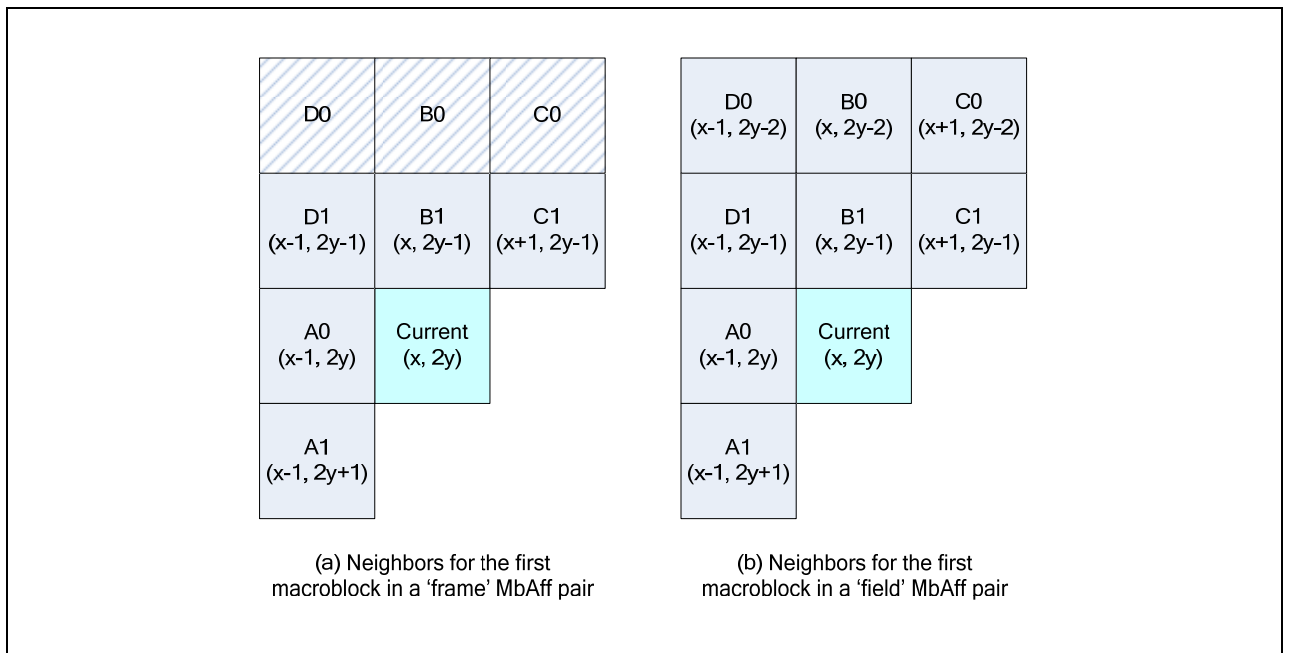
For AVD decoding, dependencies for a given macroblock may be set based on the availability of neighbor macroblocks, namely A, B, C, D and left-bottom neighbors (left-bottom only if MbAff = 1), as well as the current macroblock's address, MbAff flag and FieldMbFlag. For a macroblock in a progressive frame picture or a field picture, one macroblock may depend on up to four neighbors, A, B, C and D as shown in **Error! Reference source not found.** For a macroblock in a MbAff pair, it may depend on up to three, five or eight neighbors as shown in **Error! Reference source not found.** and **Error! Reference source not found.**, based on the current macroblock's address and FieldMbFlag.

The neighbor's availability depends on the corresponding **IntraPredAvailFlagA|B|C|D|E** flags for the macroblock (or the macroblock pair). Hardware assumes that the flags are set correctly in the MEDIA_OBJECT_EX command as shown in **Error! Reference source not found.** For simplicity, the left neighbor pair (A0 and A1) availability for a MbAff macroblock can be determined as a group by **IntraPredAvailFlagA | IntraPredAvailFlagE**. For the second macroblock in a 'frame' MbAff pair, it depends on the first macroblock in the pair and it is always available.

Neighbor addresses of a macroblock in a progressive frame picture (MbAff = 0) or a field picture with up to 4 dependencies

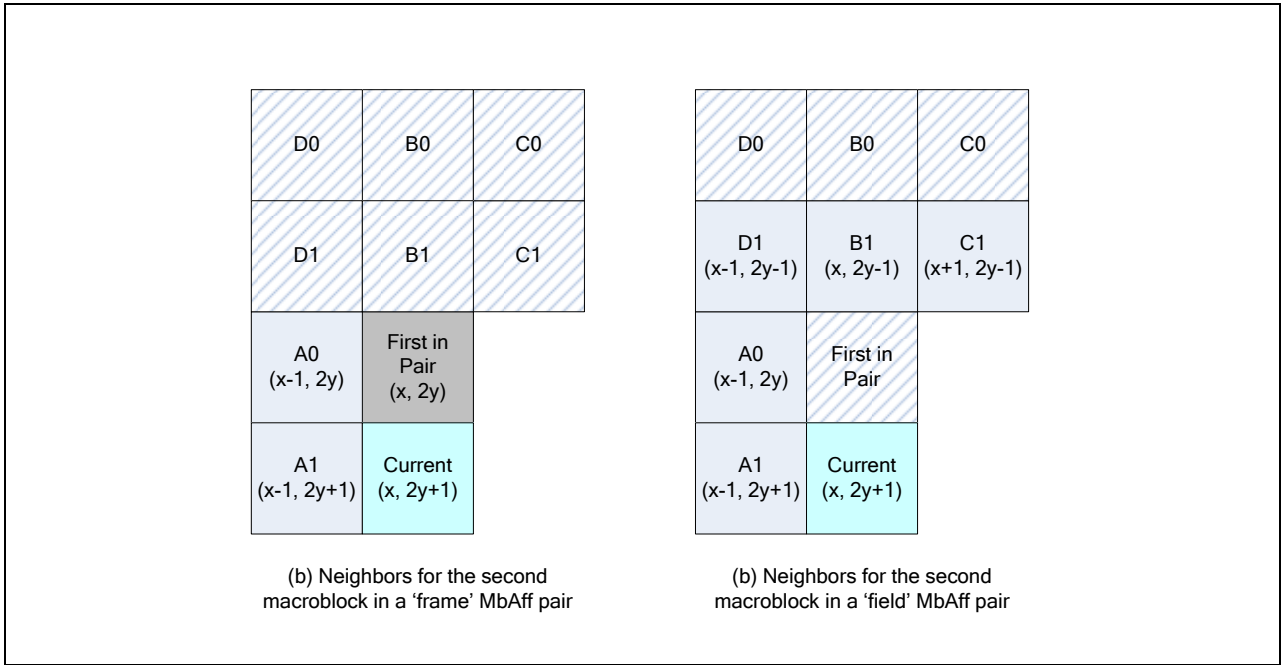


Neighbor addresses of the first macroblock in a MbAff frame picture (MbAff = 1) with up to 8 dependencies





Neighbor addresses of the second macroblock in a MbAff frame picture (MbAff = 1) with up to 8 dependencies



Neighbor Availability

| MbAff | FieldMbFlag | VertOrigin[0] | A | B | C | D | LB | Comments |
|-------|-------------|---------------|---|----|---|----|----|--|
| 0 | 0/1 | 0/1 | √ | √ | √ | √ | 0 | Progressive or Field picture |
| 1 | 0 | 0 | √ | √ | √ | √ | √ | 1 st Frame MbAff macroblock |
| 1 | 0 | 1 | √ | na | 0 | na | √ | 2 nd Frame MbAff macroblock |
| 1 | 1 | 0 | √ | √ | √ | √ | √ | 1 st Field MbAff macroblock |
| 1 | 1 | 1 | √ | √ | √ | √ | √ | 2 nd Field MbAff macroblock |



1.4.3.2 VC1-Style Dependency Example

For VC1, only one dependency may be set depending on the availability of the upper neighbor macroblock.

Macroblock sequence order in a VC-1 picture with WidthInMblk = 5 and HeightInMblk = 6

| | 0 | 1 | 2 | 3 | 4 |
|---|----|----|----|----|----|
| 0 | 0 | 1 | 2 | 3 | 4 |
| 1 | 5 | 6 | 7 | 8 | 9 |
| 2 | 10 | 11 | 12 | 13 | 14 |
| 3 | 15 | 16 | 17 | 18 | 19 |
| 4 | 20 | 21 | 22 | 23 | 24 |
| 5 | 25 | 26 | 27 | 28 | 29 |

1.4.3.3 Walker Parameter Description

The global and local loops are both described by the same four parameters:

- Resolution,
- Starting location,
- Outer unit vector,
- Inner unit vector

The local inner loop has some special modes that will be described later. A table of the user inputs and some example values are given below:

| GLOBAL LOOP PARAMETERS | | | | | | | |
|---------------------------------------|-------|--------------|---|------------------------|------------|------------------------|----|
| Global Resolution | | Global Start | | Outer Loop Unit Vector | | Inner Loop Unit Vector | |
| X | Y | X | Y | X | Y | X | Y |
| 120 | 68 | 0 | 0 | 32 | 0 | 0 | 32 |
| LOCAL LOOP PARAMETERS | | | | | | | |
| Block Resolution | | Local Start | | Outer Loop Unit Vector | | Inner Loop Unit Vector | |
| X | Y | X | Y | X | Y | X | Y |
| 32 | 32 | 0 | 0 | 1 | 0 | -2 | 2 |
| LOCAL INNER LOOP SPECIAL MODE SELECTS | | | | | | | |
| Dual Mode | Repel | Attract | | | ExtraSteps | X | Y |
| TRUE | FALSE | FALSE | | | 1 | 0 | 1 |



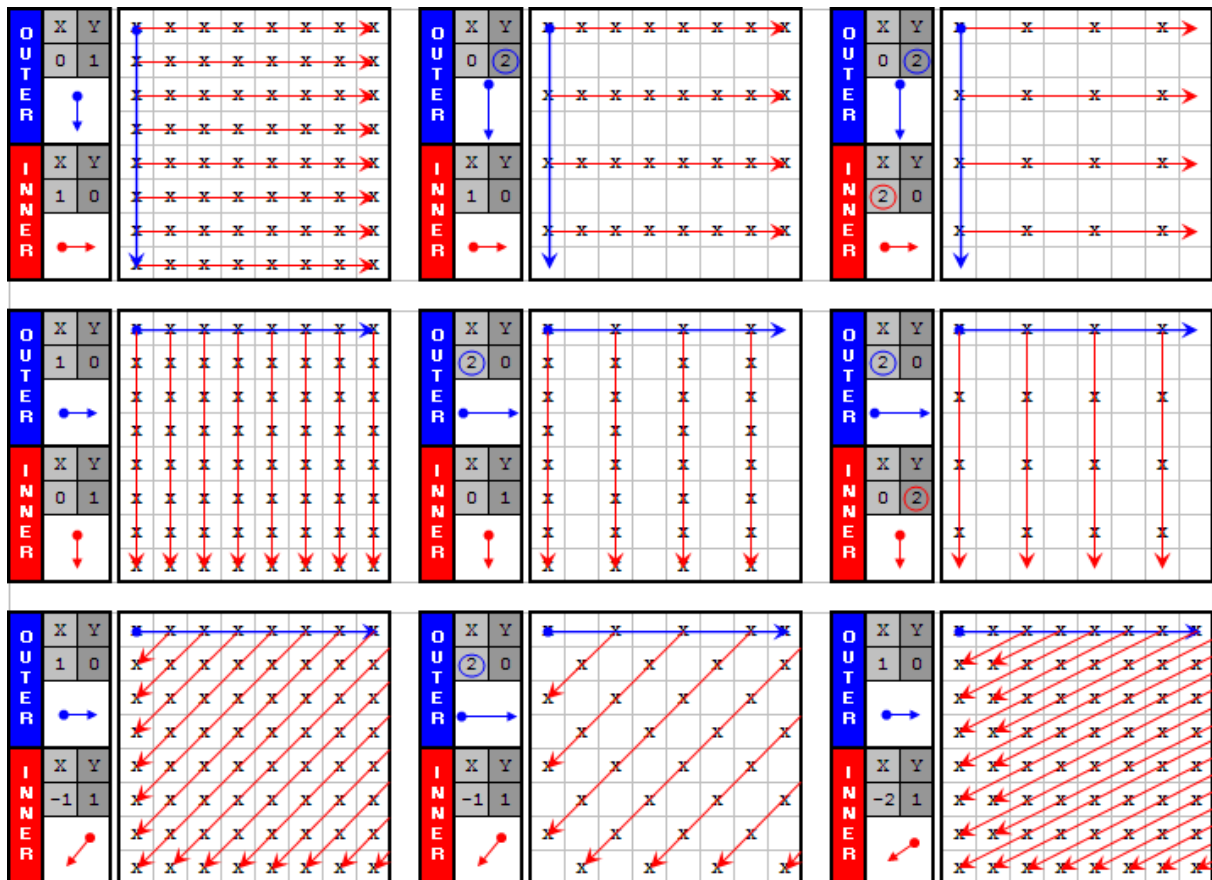
It should be emphasized that the value of what a “unit” represents is implicitly defined by the user. In other words, the walker traverses a “unit normalized space” that is not inherently bound to pixel walking. If the smallest unit of work the user wants to walk is a 4x3 block of pixels, you can program the inner loop to step (4,3) or (1,1):

- In the first case (4,3) the user is walking in units of pixels
- In the second case (1,1) the user is walking in units of 4x3 blocks of pixels.

It should be noted that hardware doesn't contain enough bits for pixel walking for pixel resolution like 1920x1088. The intended usage of the walker is for block walking whereas the block size is not relevant to the walker parameters.

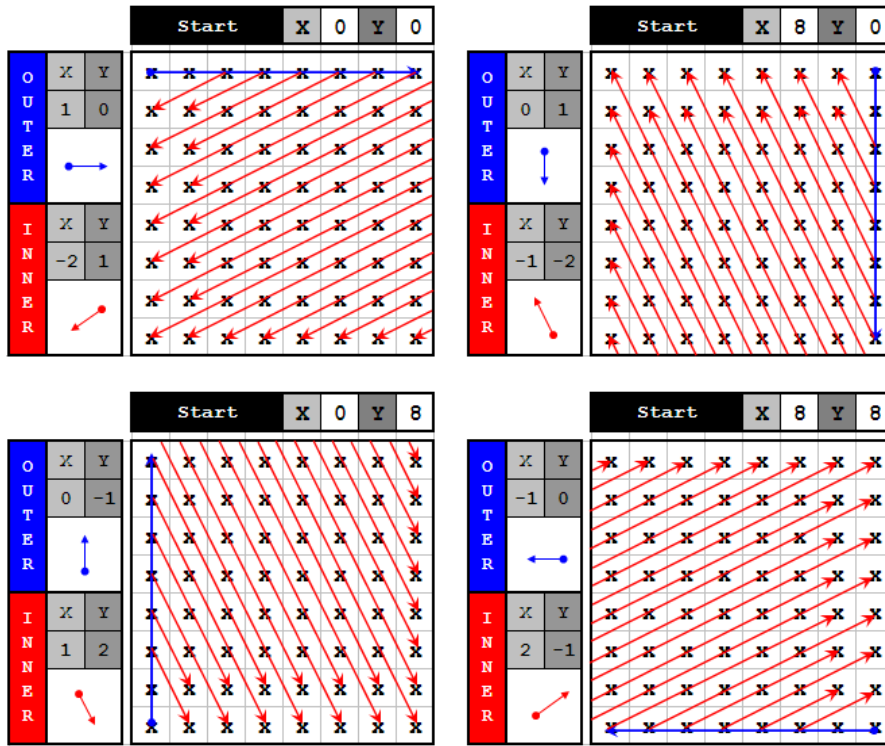
1.4.3.4 Basic Parameters for the Local Loop

The local inner and outer loop xy-pair parameters alone can describe a large variety of primitive walking patterns. Below are 9 primitive walking patterns generated by varying only the inner and outer unit step vectors of the local loop:




- The top row shows the outer unit vector pointing down (+Y) and the inner unit vector pointing right (+X). Rows and columns can easily be skipped by increasing the unit step vectors above one.
- The middle row the outer unit vector pointing right (+X) and the inner unit vector pointing down (+Y). Again, rows and columns are skipped by increasing the unit step vectors beyond one.
- The last row shows the capability to walk angles not perpendicular to the edge. The 1st shows a 45° walking pattern by setting the inner unit vector to (-1,1). The 2nd shows a checkerboard pattern by skipping every other outer loop and retaining the inner unit vector of (-1,1). The 3rd shows a 26.5° walking pattern by setting the inner unit vector to (-2,1).

The block resolution, shown as [8,8], and the starting location, currently [0,0], can be varied and the above patterns can be stretched and rotated many ways. The diagram below shows an example of where the start position and unit step vectors can be set to achieve a full rotation of the same pattern:




1.4.3.5 Dual Mode of Local Loop

The local Inner Loop Special mode selects are included to aid in the distribution of work, specifically with two slices in mind. Essentially, the local inner loop can be bisected and each half-walk can be directed inward towards the center of the image (dual). The local inner loop need not be bisected, and can either move away from the outer loop (repel) or move towards it (attract) when an even split is not desired:

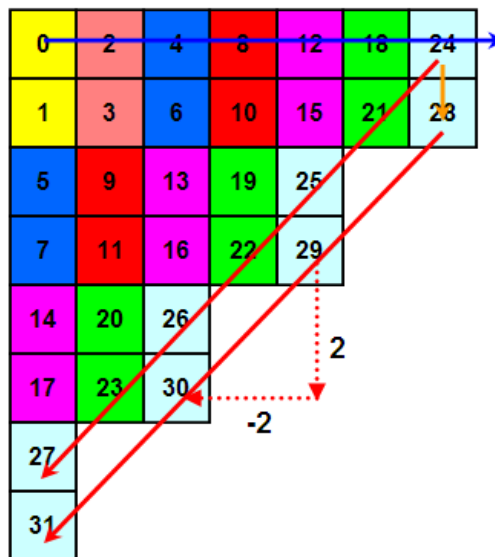


In Dual mode, the sequence will alternate between two half-walks such that every-other output would go to the same slice. This effect will produce a more balanced workload to two slices as shown in the example below where the color of the block represents which slice it was dispatched to. This is the walker's approach to fine-grained parallelism.



1.4.3.6 MbAff-Like Special Case in Local Loop

The local loop has an additional middle loop that is used to achieve some specific walking patterns, with MBAFF mode especially in mind. A pattern to handle MBAFF AVC content is to walk the top macroblocks of all macroblock pairs (MB-pairs) on a wavefront followed by the respective bottom macroblocks. The pattern is shown below.



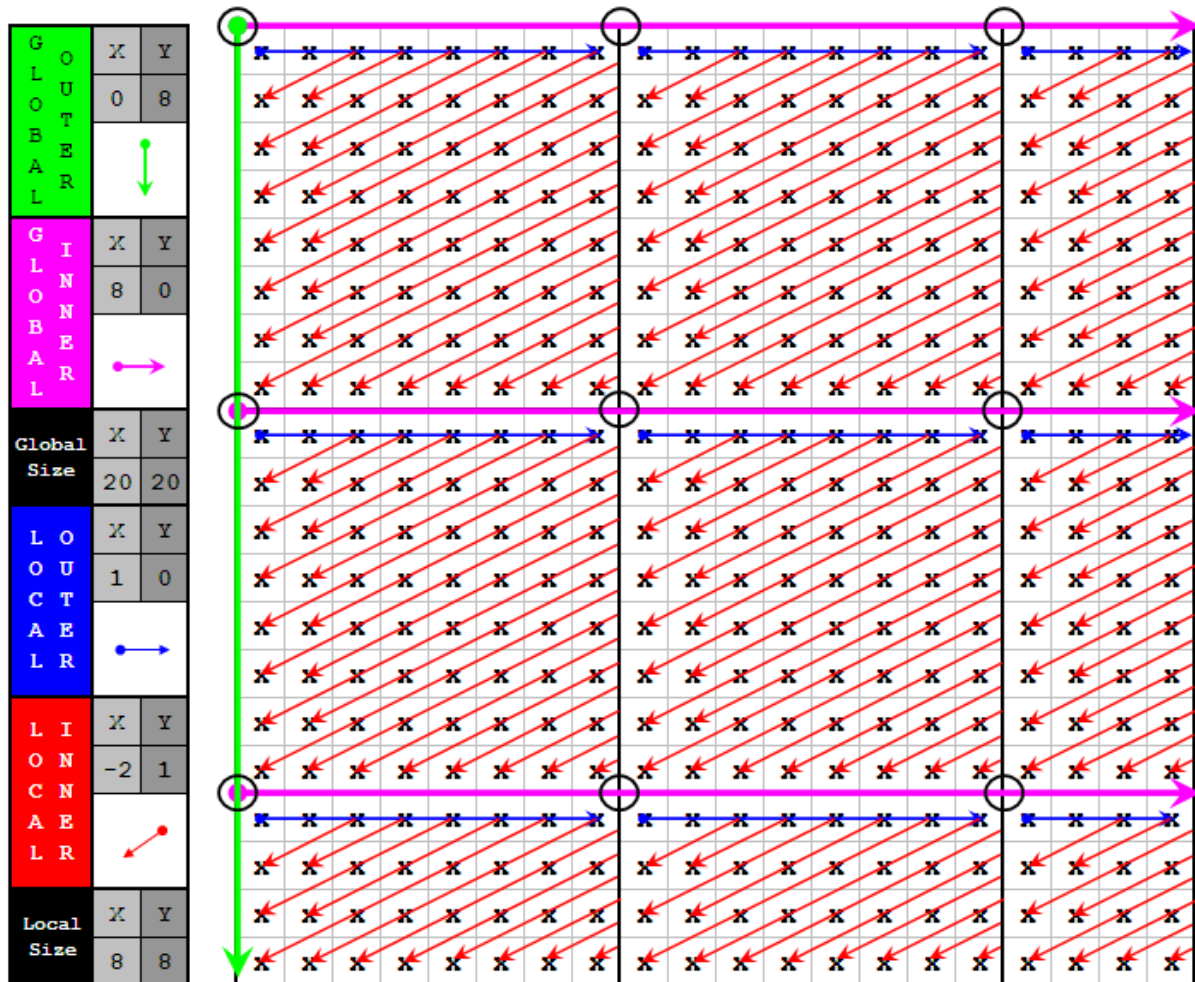
The outer loop unit step vector would be $[1, 0]$ and the inner loop unit step vector would be $[-2, 2]$. A third loop is necessary to repeat the inner loop, only shifted down a unit before restarting. Thus, a middle loop with a unit step vector of $[0, 1]$ would achieve this MBAFF pattern. Additionally, the number of “extra steps” taken by the middle loop would be 1 in this case.



The addition of a middle loop also creates more overall flexibility, which seems necessary due to the integer-based unit step vector solution proposed (Manhattan distance issues etc.).

1.4.3.7 Global Loop

The same set of general parameters is used to describe the global loop as well. Thus, a global loop that is walking a raster-scan pattern can be combined with a local loop that is walking a 26.5° pattern (or vice-versa). As shown in the example below, if the local block size ([8,8]) is not an even multiple of the global resolution ([20,20]), the slack is still processed by dynamically changing the local block resolution.



The global loop will always resolve to be the upper-left corner of the local loop, shown above black circles. Note that local loop can still start in any corner of the local block, but the local (0,0) will always be the location where global loop begins the local loop, hence the upper-left corner.

The user can specify where the starting location of the global loop as with the local loop. If the user were to set the global starting location to (16,16) in the previous example, after inverting the global outer and global inner unit step vectors the same pattern would be achieved in the reverse order. Note that the slack would still be handled along the right and bottom edge of the global image in that case. The user could have also started at (12,12) in which case the slack would be handled on the left and top faces.



Walker Algorithm Description

The walker algorithm has been tested and optimized in software. A high-level pseudo-code description is given below:

```
Walker(){ //C-Style Pseudo-Code of Walker Algorithm
    Load_Inputs_And_Initialize();
    While (Global_Outer_Loop_In_Bounds()){
        Global_Inner_Loop_Intialization();
        While (Global_Inner_Loop_In_Bounds()){
            Local_Block_Boundary_Adjustment();
            Local_Outer_Loop_Initialization();
            While (Local_Outer_Loop_In_Bounds()){
                Local_Middle_Loop_Initialization();
                While (Local_Middle_Steps_Remaining()){
                    Local_Inner_Loop_Initialization();
                    While (Local_Inner_Loop_Is_Shrinking()){
                        Execute();
                        Calculate_Next_Local_Inner_X_Y();
                    } //End Local Inner Loop
                    Calculate_Next_Local_Middle_X_Y();
                } //End Local Middle Loop
                Calculate_Next_Local_Outer_X_Y();
                Calculate_Next_Local_Inverse_Outer_X_Y();
            } //End Local Outer Loop
            Calculate_Next_Global_Inner_X_Y();
        } //End Global Inner Loop
        Calculate_Next_Global_Outer_X_Y();
    } //End Global Outer Loop
} //End Walker
```



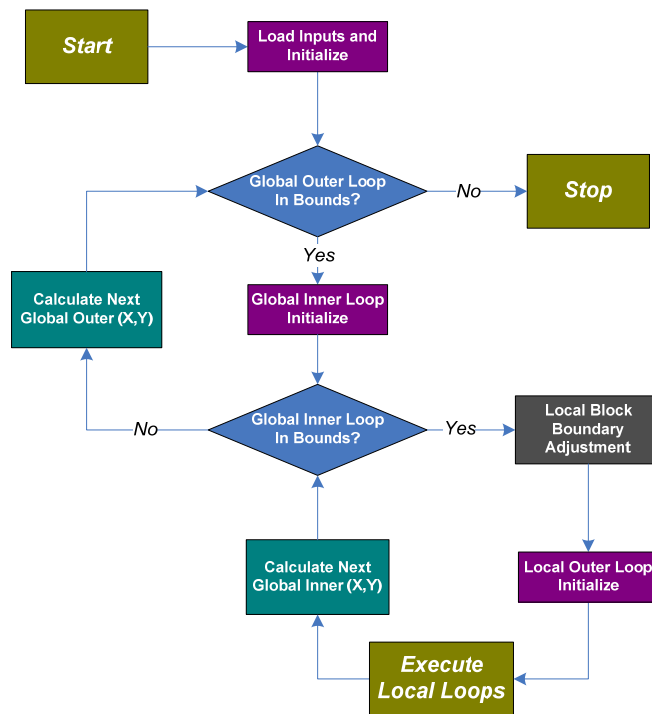
The pseudo-code has the following characteristics:

- There are 5 levels of iteration
- The highest 2 levels are called “global” and the lowest 3 levels are called “local”
 - The global loop is split into an outer and an inner loop.
 - The local loop is split into an outer, a middle, and an inner loop.
 - A bounding box for the global and local resolution is defined by the user.
 - The starting location within each bounding box is also specified by the user.
- Each of the 5 loops has its own persistent
 - Current position (x,y)
 - Unit step vector (x,y)
- The final output (x,y) is a summation of the global x,y and the local x,y.
- The next (x,y) for given level can be calculated while the next lower level is still executing. Additionally, the result can be used to check to see if the current level will execute again once control is returned.

The flow of the global outer and inner loops is:

1. Check a bound condition
2. Initialize the next level loop
3. Execute the next level loop
4. When the next level loop fails its condition, calculate the next position for the current loop level and repeat.

Figure 1-2. Walker algorithm flowchart for the Global Loop

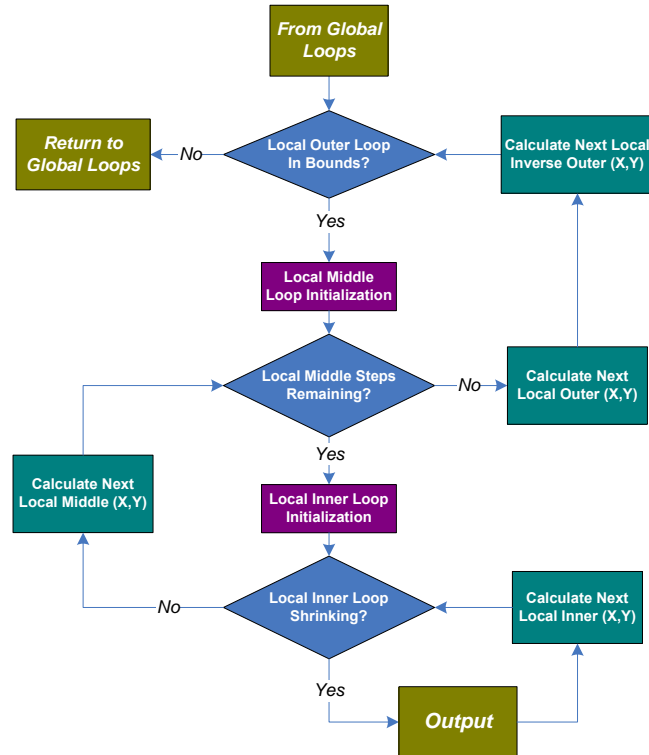


Take note of the grey box “Local Block Boundary Adjustment”. This logic is necessary to adjust the local block size when the distance between the current global position to the edge of the image is less than the local resolution. Additionally, the local starting positions might be modified here as well if the defined starting position is larger than the new local block size.

The flow of the 3 local loops does not vary much from the 2 global loops. The differences are:

- In addition to a boundary check, the local middle loop also ensures the number of middle steps is less than or equal to the user defined “number of extra steps”.
- The local inner loop only checks to see if the prior distance between the x,y starting and ending points are greater than their current distance. If this is true, it implies that the two inner loops are converging towards each other.
- When the middle loop check fails, both the starting points (local outer) and ending points (local inner) are updated.

Figure 1-3. Walker algorithm flowchart for the Local Loop

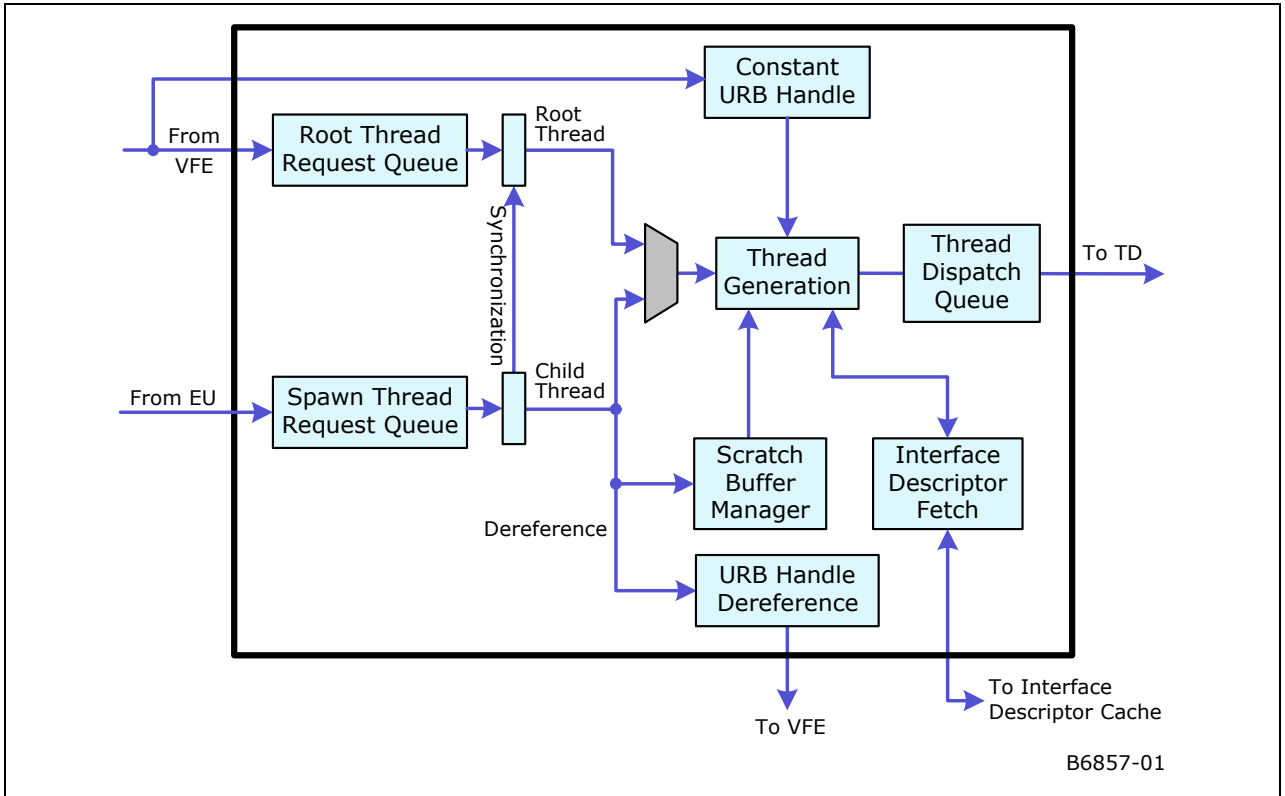


1.5 Thread Spawner Unit

The Thread Spawner (TS) unit is responsible for making thread requests (root and child) to the Thread Dispatcher, managing scratch memory, maintaining outstanding root thread counts, and monitoring the termination of threads.



Thread Spawner block diagram



1.5.1 Basic Functions

1.5.1.1 Root Threads Lifecycle

Thread requests sourced from VFE are called **root threads**, since these threads may be creating subsequent (child) threads. A root thread may be a macroblock thread created by VFE as in VLD mode, or may be a general-purpose thread assembled by VFE according to full description provided by host software in Generic mode.

Thread requests are stored in the Root Thread Queue. TS keeps everything needed to get the root threads ready for dispatch and then tracks dispatched threads until their retirement.

TS arbitrates between root thread and child thread. The root thread request queue is in the arbitration only if the number of outstanding threads does not exceed the maximum root thread state variable. Otherwise, the root thread request queue is stalled until some other root threads retire/terminate.



Once a root thread is selected to be dispatched, its lifecycle can be described by the following steps:

1. TS forwards the interface descriptor pointer to the L1 interface descriptor cache (a small fully associated cache containing up to 4 interface descriptors). The interface descriptor is either found in the cache or a corresponding request is forwarded to the L2 cache. Interface descriptors return back to TS in requesting order.
 - Once TS receives the interface descriptor, it checks whether maximum concurrent root thread number has reached to determine whether to make a thread dispatch request or to stall the request until some other root threads retire. If the thread requests the use of scratch memory, it also generates a pointer into the scratch space.
2. TS then builds the transparent header and the R0 header.
3. Finally, TS makes a thread request to the Thread Dispatcher.
4. TS keeps track of dispatched thread, and monitors messages from the thread (resource dereference and/or thread termination). When it receives a root thread termination message, it can recover the scratch space and thread slot allocated to it. The URB handle may also be dereferenced for a terminated root thread for future reuse. It should be noted that URB handle dereference may occur before a root thread terminates. See detailed description in the Media Message section.
 - It is the root thread's responsibility (software) to guarantee that all its children have retired before the root thread can retire.

1.5.1.2 URB Handles

VFE is in charge of allocating URB handles for root threads. One URB handle is assigned to each root thread. The handle is used for the payload into the root thread.

If Children Present state variable is not set (root-without-child mode), TS signals VFE to dereference the URB handle immediately after it receives acknowledgement from TD that the thread is dispatched.

If Children Present state variable is set (root-with-child mode), the URB handle is forwarded to the root thread and serves as the return URB handle for the root thread. TS does not signal dereference at the time of dispatch. TS signals URB handle dereference only when it receives a resource dereference message from the thread.

[Pre-SNB] Children Present is a state variable

1.5.1.3 Root to Child Responsibilities

Any thread created by another thread running in an EU is called a **child thread**. Child threads can create additional threads, all under the tree of a root which was requested via the VFE path.

A root thread is responsible of managing pre-allocated resources such as URB space and scratch space for its direct and indirect child threads. For example, a root thread may split its URB space into sections. It can use one section for delivering payload to one child thread as well as forwarding the section to the child thread to be used as return URB space. The child thread may further subdivide the URB section into subsections and use these subsections for its own child threads. Such process may be iterated. Similarly, a root thread may split its scratch memory space into sections and give one scratch section for one child thread.

TS unit only enforces limitation on number of outstanding root threads. It is the root threads' responsibility to limit the number of child threads in their respected trees to balance performance and avoid deadlock.

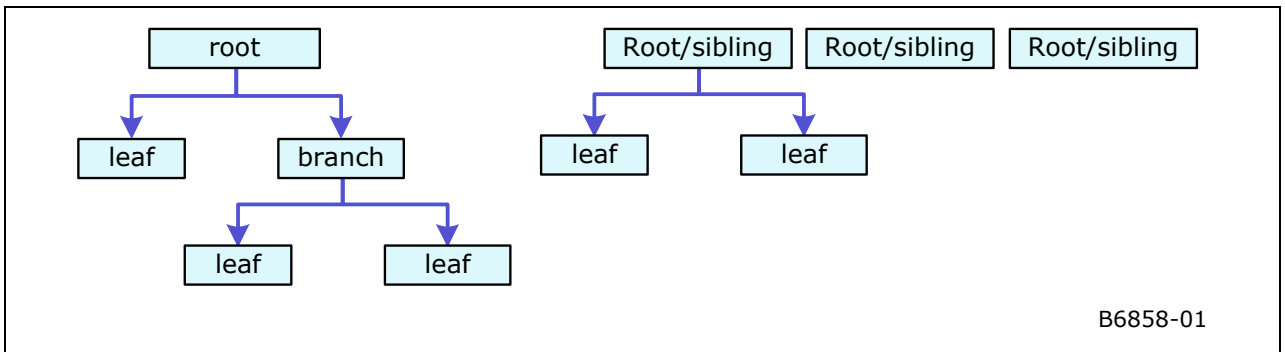


1.5.1.4 Multiple Simultaneous Roots

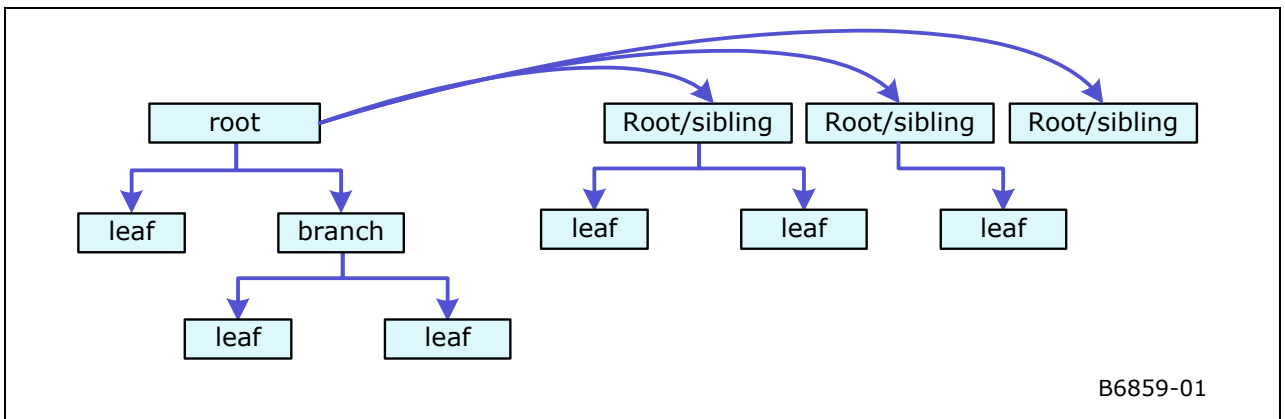
Multiple root threads are allowed concurrently running in GEN4 execution units. As there is only one scratch space state variable shared for all root threads, all concurrent root thread requiring scratch space share the same scratch memory size. **Error! Reference source not found.** depicts two examples of thread-thread relationship. The left graph shows one single tree structure. This tree starts with a single root thread that generates many child threads. Some child threads may create subsequent child threads. The right graph shows a case with multiple disconnected trees. It has multiple root threads, showing sibling roots of disconnected trees. Some roots may have child threads (branches and leafs) and some may not.

There is another case (as shown in **Error! Reference source not found.**) where multiple trees may be connected. If a root is a synchronized root thread, it may be dependent on a preceding sibling root thread or on a child thread.

Examples of thread relationship



A example of thread relationship with root sibling dependency





1.5.1.5 Synchronized Root Threads

A synchronized root thread (SRT) originates from a MEDIA_OBJECT command with Thread Synchronization field set. Synchronized root threads share the same root thread request queue with the non-synchronized roots. A SRT is not automatically dispatched. Instead, it stays in the root thread request queue until a spawn-root message is at the head of the child thread request queue. Conversely, a spawn-root message in the child thread request queue will block the child thread request queue until the head of root thread request queue is a SRT. When they are both at the head of queues, they are taken out from the queue at the same time.

A spawn-root message may be issued by a root thread or a child thread. There is no restriction. However, the number of spawn-root messages and the number of SRT must be identical between state changes. Otherwise, there can be a deadlock. Furthermore, as both requests are blocking, synchronized root threads must be used carefully to avoid deadlock.

When Scoreboard Control is enabled, the dispatch of a SRT originated from a MEDIA_OBJECT_EX command is still managed by the same way in addition to the hardware scoreboard control.

1.5.1.6 Deadlock Prevention

Root threads must control deadlock within their own child set. Each root is given a set of preallocated URB space; to prevent deadlock it must make sure that all the URB space is not allocated to intermediate children who must create more children before they can exit.

There are limits to the number of concurrent threads. The upper bound is determined by the number of execution units and the number of threads per EU. The actual upper bound on number of concurrent threads may be smaller if the GRF requirement is large. Deadlock may occur if a root or intermediate parent cannot exit until it has started its children but there is no space (for example, available thread slot in execution units) for its children to start.

To prevent deadlock, the maximum number of root threads is provided in VFE state. The Thread Spawner keeps track of how many roots have been spawned and prevents new roots if the maximum has been reached. When child threads are present, it is software's responsible of constraining child thread generation, particularly the generation of child threads that may also spawn more child threads.

Child thread dispatch queue in TS is another resource that needs to be considered in preventing deadlock. The child thread dispatch queue in TS is used for (1) message to spawn a child thread, (2) message to spawn a synchronized root thread, and (3) thread termination message. If this queue is full, it will prevent any thread to terminate, causing deadlock.

For example, if an application only has one root thread (max # of root threads is programmed to be one). This root thread spawns child threads. In order to avoid deadlock, the maximum number of outstanding child thread that this root thread can spawn is the sum of the maximum available thread slots plus the depth of the child thread dispatch queue minus one.

$$\text{Max_Outstanding_Child_Threads} = (\text{Thread Slot Number} - 1) + (\text{TS Child Queue Depth} - 1)$$

Adding other root threads (synchronized and/or non-synchronized) to the above example, the situation is more complicated. A conservative measure may have to use to prevent deadlock. For example, the root thread spawning child threads may have to exclude the max number of root threads as in the following equation to compute the maximum number of outstanding child threads to be dispatched.



$\text{Max_Outstanding_Child_Threads} = (\text{Thread Slot Number} - 1) + (\text{TS Child Queue Depth} - 1) - (\text{Max Root Threads} - 1)$

TS Resource Available in Device Hardware

| Device | Child Thread Dispatch Queue Depth |
|----------|-----------------------------------|
| [DevBW] | 8 |
| [DevCL] | 8 |
| [DevCTG] | 8 |
| [DevILK] | OPEN |

1.5.1.7 Child Thread Lifecycle

When a (parent) thread creates a child thread, the parent thread behaves like a fixed function. It provides all necessary information to start the child thread, by assembling the payload in URB (including R0 header) and then sending a spawn thread message to TS with following data:

- An interface descriptor pointer for the child thread.
- A pointer for URB data

The interface descriptor for a child may be different from the parent – how the parent determines the child interface descriptor is up to the parent, but it must be one from the interface descriptor array on the same interface descriptor base address.

The URB pointer is not the same as a URB handle. It does not have an URB handle number and does not appear in any handle table. This is acceptable because the URB space is never reclaimed by TS after a child is dispatched, but rather when the parent releases its original handles and/or retires.

The R0 header for a child, as part of the URB payload, consisting of the 32-bit field from the parent and a parent created field to uniquely identify the child.

The child request is stored in the child thread queue. The depth of the queue is limited to 8, overrun is prevented by the message bus arbiter which controls the message bus. The arbiter knows the depth of the queue and will only allow 8 requests to be outstanding until the TS signals an entry has been removed.

As mentioned previously, child threads have higher priority over root threads. Once TS selects a child thread to dispatch, it follows these steps:

1. TS forwards the interface descriptor pointer to the L1 interface descriptor cache (a small fully associated cache containing up to 4 interface descriptors). The interface descriptor is either found in the cache or a corresponding request is forwarded to the L2 cache. Interface descriptors return back to TS in requesting order.
2. TS then builds the transparent header but not the R0 header.
3. Finally, TS makes a thread request to the Thread Dispatcher.
4. Once the dispatch is done, TS can forget the child – unlike roots, no bookkeeping is done that has to be updated when the child retires.



If more data needs to be transferred between a parent thread and its child thread than that can fit in a single URB payload, extra data must be communicated via shared memory through data port.

1.5.1.8 Arbitration between Root and Child Threads

When both root thread queue and child thread queue are both non-empty, TS serves the child thread queue. In other words, child threads have higher priority over root threads. The only condition that the child thread queue is stalled by the root thread queue is that the head of child thread queue is a root-synchronization message and the head of root thread queue is not a synchronized root thread.

1.5.1.9 Persistent Root Thread [DevCTG+]

Persistent Root Thread (PRT) is a persistent root thread in general stays in the system for a long period of time. It is normally a parent thread, and only one PRT is allowed in the system at a time. Upon context switch interrupt, instead of proceeding to completion, a PRT can save its software context and terminate. The PRT can be restarted later, *even if it had completed normally the last time it was executed*. Therefore, the PRT must always save enough context (via data port messages to a predefined surface) to allow it to restart from where it left off (including determining that it has nothing left to do). However, since only one PRT can execute at a time, once the next PRT starts, the previous one will never be restarted, thus the context save surface can be reused from one PRT to the next.

A PRT may check the Thread Restart Enable bit in the R0 header to find out whether it is a fresh start or resumed from a previous interrupt and then can continue operations from that previously saved context.

PRT can be interleaved with other root (such as parent root thread, or synchronized root thread) and child threads. A parent root thread is not necessarily a PRT, and doesn't have to be as long as it can be finished in deterministic time that is shorter than required for fine-grain context switch interrupt.

Use of PRT must follow the following rules:

- There can only be one PRT in the media pipeline at a given time. That means, there shall not be any other media primitive commands (MEDIA_OBJECT or MEDIA_OBJECT_EX) between it and the previous MI_FLUSH command. In other words, when multiple such PRTs are used in a sequence of media primitive commands, MI_FLUSH must be inserted.

1.5.1.10 GPGPU Functions

In GPGPU mode, the Thread Spawner allocate and track the barriers and shared local memory per thread group.

1.5.1.10.1 Thread Group Tracking

The TSG needs to keep track of the threads outstanding in a group to know when the thread group barrier and Shared Local Memory can be reclaimed. This can be done by keeping a counter per active thread group (up to 16 per half-slice) which increments when a new thread is sent out and decremented when the thread retires. The assigned barrier ID (with half-slice bit) is unique per thread group and much smaller than the thread group id and so will be used to keep track of the thread group instead.



Since TSL sends the thread retirement via the Message Channel rather than the thread retirement bus, the barrier ID used to identify the thread group can be sent at the same time. A CAM will then match the ID with the counter to decrement.

There is a potential corner case of a thread group without barriers being partly dispatched, then retiring before the rest of the thread group is sent. This should be OK, since the lack of barriers means that there are no dependencies between threads.

1.5.1.10.2 Shared Local Memory Allocation

The Shared Local Memory is a 64k block per half-slice in the L3 that must be shared between all thread groups on that half-slice. A new memory manager similar to the Scratch Space memory manager is used to allocate this space.

We are only dispatching threads from a single Interface Descriptor at a time. If a new Interface Descriptor is requested the pipe is drained and all shared memory recovered before starting to allocate new shared memory. This means that only a single size of shared memory needs to be supported at once.

For simplicity, only power-of-2 sizes from 4k to 64k are allowed. The thread request will specify how much is needed. The first thread of a Thread Group is marked as requiring a new shared local memory – if not the old Shared Local Memory offset is sent with the dispatch.

A simple set of 16-bits is used to allocate 4k shared memory, with fewer bits used for larger sizes. A priority encoder finds the first unused bit and the offset remembered as being associated with a particular barrier id. The barrier id is then used to track the thread group.

When the Thread Group Tracking indicates that a thread group is completely retired, that section of shared local memory can be reclaimed.

1.5.1.10.3 Software Managed Shared Local Memory

Software can optionally manage shared local memory. In this case, each thread command or thread group command will have the shared memory offset included – each command in a thread group must have the same offset, of course. If the offset requested is still being used then the command is stalled until the thread group using that offset is done.

Hardware will track the usage of this section of shared memory as before, recording the offset as being used and recording it as being available after the thread group is done.

1.5.1.10.4 Automatic Barrier Management

Instead of the barrier id in the Interface Descriptor, there is now a thread count per thread group. If a new thread group id comes in without a barrier allocated (checked with a CAM match across



16 barriers), the TSG picks a unused barrier and sends this count in a message to GWunit. It then needs to wait for an accept message back from GW before sending the dispatch to ensure that a barrier message doesn't arrive at the GW before the barrier is programmed. The barrier id picked is sent with every dispatch from this thread group.

When the thread group tracker determines that a thread group has finished, the barrier becomes available to new thread groups.

1.5.1.10.5 Local Memory / Scratch Space

The Local Memory (not to be confused with Shared Local Memory, which is shared by all thread in a thread group) is allocated per thread dispatched to the EU. The existing Scratch Space manager is used to provide between 1k and 12k bytes memory per thread.

1.5.1.10.6 Dispatch Payload

The payload for a general purpose thread will have to include the execution mask with a bit per 32-channel. SIMD16 and SIMD8 use the LSB bits of the execution mask. The 5-bit number transferred from VFE will be expanded to produce the 32-bit mask. This will use the Dmask currently used by the pixel shader dispatch in the transparent header.

1.5.2 Interfaces

1.5.2.1 Interface to VFE

TS receives an interface descriptor pointer and a URB handle from VFE. It uses the interface descriptor pointer to fetch the interface descriptor. TS uses the information in the interface descriptor along with the URB handle to fill out the transparent header in the message to TD for all threads. For root thread, TS also generate the R0 header.

TS transmits URB handle dereference signal to VFE. As described previously, the dereference signal may be at dispatch time or at later time depending on Children Present state variable. No matter which case, there is one and only one URB handle dereference for a thread.

1.5.2.2 Interface to Thread Dispatcher

TS creates the transparent header, assembles the URB handles and calls TD to dispatch a new thread. For an unsynchronized root thread, there is one URB handle managed by VFE and optionally one Constant URB handle managed by CS. For a synchronized root thread, there is one URB handle managed by VFE, a URB handle created by the synchronizing thread (the one that sends the 'spawn root thread' message, and optionally one Constant URB handle managed by CS. For a child thread, there is one URB handle managed by the parent thread plus an optional Constant URB handle.



1.6 Media State ([Pre-DevSNB])

1.6.1 Media State Model

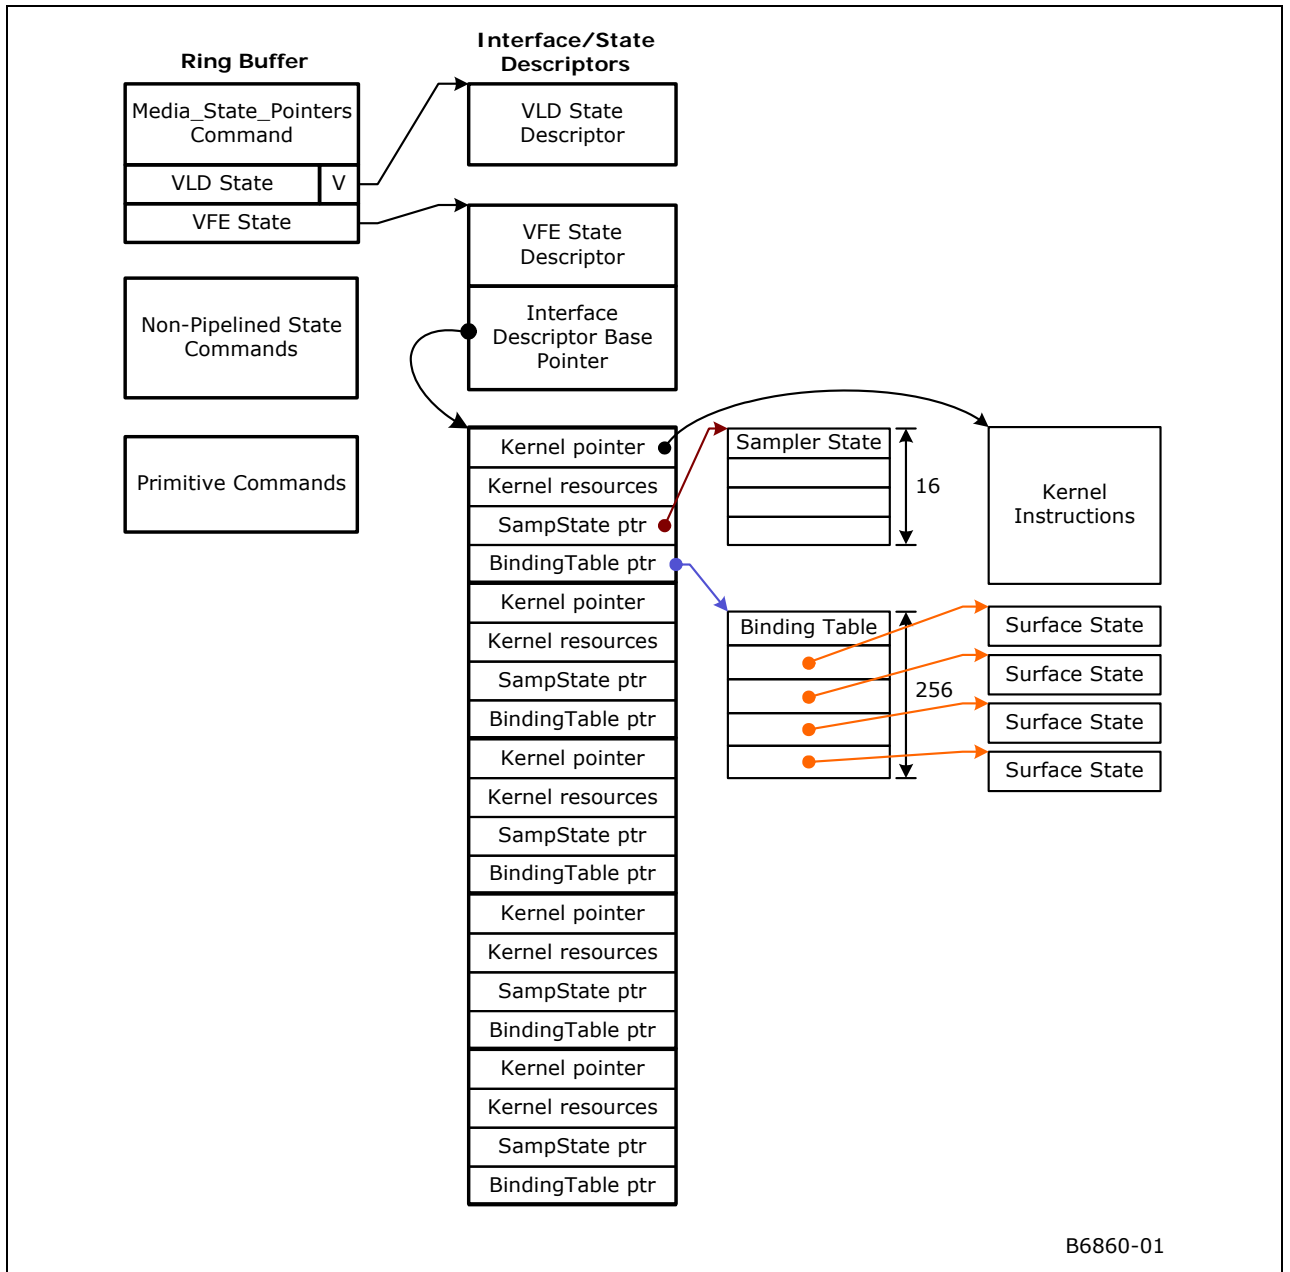
The media state model is based on an indirect state fetching mechanism. State Descriptors provide state information for fixed function units of the media pipeline. Interface Descriptors provide state information for kernels (threads) dispatched from the media pipeline. They are organized in different memory locations.

VFE State Descriptor contains states for both VFE unit and TS unit. The special purpose VLD state information is provided by a separate VLD State Descriptor.

All Interface Descriptors have the same size and are organized as a contiguous array in memory. They can be selected by Interface Descriptor Index for a given kernel. This allows different kinds of kernels to coexist in the system.

The MEDIA_STATE_POINTERS command provides the memory pointers to the Descriptors.

Media State Model





1.6.2 VFE_STATE

| Dword | Bit | Description |
|-------|-------|---|
| 0 | 31:10 | Scratch Space Base Pointer. Specifies the 1k-byte aligned address offset to scratch space for use by the kernel. This pointer is relative to the General State Base Address . Format = GeneralStateOffset[31:10] |
| | 9:8 | Reserved: MBZ |
| | 7 | Extended VFE State Present. This field specifies whether extended VFE state is present or not. It must be programmed with the same value as the Extended VFE State Enable field in MEDIA_STATE_POINTERS command. 0 = Disabled. No extended VFE state (and Extension State Pointer is ignored). 1 = Enabled. The extended VFE state pointed by Extended State Pointer is loaded [DevBW, DevCL] This field is reserved and MBZ. |
| | 6:4 | Reserved : MBZ |
| | 3:0 | Per Thread Scratch Space. Specifies the amount of scratch space allowed to be used by each thread. The driver must allocate enough contiguous scratch space, pointed to by the Scratch Space Pointer, to ensure that the Maximum Number of Threads each get Per Thread Scratch Space size without exceeding the driver-allocated scratch space. Note: The definition of this field is different from that in 3D fixed functions, where the per-thread scratch space is specified in powers of 2. Format = U4 Range = [0,11] indicating [1k bytes, 12k bytes] |
| 1 | 31:25 | Maximum Number of Threads. Specifies the maximum number of simultaneous root threads allowed to be active. Used to avoid using up the scratch space, or to avoid potential deadlock. Note that MSB will be zero due to the range limit below. [Pre-DevSNB]: Format = U7 representing (thread count – 1) Range = [0, n-1] where n = (# EUs) * (# threads/EU). See <i>Graphics Processing Engine</i> for listing of #EUs and #threads in each device. |
| | 24:16 | URB Entry Allocation Size [Pre-DevSNB]. Specifies the length of each URB entry used by the unit, in 512-bit register increments - 1. Format = U9 Range = [0,255] indicating [1,256] 512-bit register increments |
| | 15:9 | Number of URB Entries [Pre-DevSNB]. Specifies the number of URB entries that are used by the unit. Format = U7 Range = [1,64] |
| | 8:7 | Reserved : MBZ |



| Dword | Bit | Description |
|-------|------|--|
| | 6:3 | VFE Mode 0000 – Generic Mode 0001 – VLD Mode (MPEG-2 only) 0010 – IS Mode (supporting WMV IDCT) 0100 – AVC-MC Mode 0111 – AVC-IT Mode 1011 – VC1-IT Mode All other encodings are reserved [DevBW, DevCL] AVC-MC, AVC-IT and VC1-IT modes are not supported [DevBW] VLD mode is not supported |
| | 2 | Children Present. Indicates that the root thread may send spawn messages to spawn child threads and/or synchronized root threads. In VLD Mode, this field must be 0. Format = Enable |
| 2 | 31:4 | Interface Descriptor Base Pointer. Specifies the 16-byte aligned address of the interface descriptor base pointer. This pointer is relative to the General State Base Address . Format = GeneralStateOffset[31:4] |
| | 3:0 | Reserved : MBZ |

1.6.3VLD_STATE

| Dword | Bit | Description |
|-------|-------|--|
| 0 | 31:28 | f_code[1][1]. Used for backward motion vector prediction. See ISO/IEC 13818-2 §7.6.3.1 for details |
| | 27:24 | f_code[1][0]. Used for backward motion vector prediction. See ISO/IEC 13818-2 §7.6.3.1 for details |
| | 23:20 | f_code[0][1]. Used for forward motion vector prediction. See ISO/IEC 13818-2 §7.6.3.1 for details |
| | 19:16 | f_code[0][0]. Used for forward motion vector prediction. See ISO/IEC 13818-2 §7.6.3.1 for details |
| | 15:14 | Intra DC Precision. See ISO/IEC 13818-2 §6.3.10 for details. |
| | 13:12 | Picture Structure. This field specifies whether the picture is encoded in the form of a frame picture or one field (top or bottom) picture. See <i>ISO/IEC 13818-2</i> §6.3.10 for details. Format = MPEG_PICTURE_STRUCTURE 00 = Reserved 01 = MPEG_TOP_FIELD 10 = MPEG_BOTTOM_FIELD 11 = MPEG_FRAME |



| Dword | Bit | Description | | | | | | | | | |
|-------|--|--|----------------------------------|-------------------------------|----------------------------------|------------------|---------|---------|------------------|---------|---------|
| | 11 | <p>TFF (Top Field First). When two fields are stored in a picture, this bit indicates if the top field is the first field.</p> <p>For a frame P picture, the value 1 indicates that the top field of the reconstructed frame is the first field output by the decoding process, the same as defined in ISO/IEC 13818-2 §6.3.10. Particularly, it is used by the hardware to calculate derivative motion vectors from the dual-prime motion vectors.</p> <p>For a field P picture, hardware uses this bit together with the Picture Structure to determine if the current picture is the Second Field. In this case, the definition of this bit differs from ISO/IEC 13818-2 §6.3.10 – software must derive the value for this bit according to the following relation:</p> <table border="1" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th></th> <th>Picture Structure = top field</th> <th>Picture Structure = bottom field</th> </tr> </thead> <tbody> <tr> <td>Second Field = 0</td> <td>TFF = 1</td> <td>TFF = 0</td> </tr> <tr> <td>Second Field = 1</td> <td>TFF = 0</td> <td>TFF = 1</td> </tr> </tbody> </table> | | Picture Structure = top field | Picture Structure = bottom field | Second Field = 0 | TFF = 1 | TFF = 0 | Second Field = 1 | TFF = 0 | TFF = 1 |
| | | Picture Structure = top field | Picture Structure = bottom field | | | | | | | | |
| | Second Field = 0 | TFF = 1 | TFF = 0 | | | | | | | | |
| | Second Field = 1 | TFF = 0 | TFF = 1 | | | | | | | | |
| | 10 | Frame Prediction Frame DCT. This field provides constraints on the DCT type and prediction type. It affects the syntax of the bitstream. | | | | | | | | | |
| | 9 | Concealment Motion Vector Flag. This field indicates if the concealment motion vectors are coded in intra macroblocks. It affects the syntax of the bitstream. | | | | | | | | | |
| | 8 | <p>Quantizer Scale Type. This field specifies the quantizer scaling type.</p> <p>Format = MPEG_Q_SCALE_TYPE 0: MPEG_QSCALE_LINEAR 1: MPEG_QSCALE_NONLINEAR</p> | | | | | | | | | |
| 7 | Intra VLC Format. This field is used by VLD. | | | | | | | | | | |
| 6 | <p>Scan Order. This field specifies the Inverse Scan method for the DCT-domain coefficients in the blocks of the current picture.</p> <p>Format = MPEG_INVERSESCAN_TYPE 0 = MPEG_ZIGZAG_SCAN 1 = MPEG_ALTERNATE_VERTICAL_SCAN</p> | | | | | | | | | | |
| 5:0 | Reserved. | | | | | | | | | | |
| 1 | 31:14 | Reserved. | | | | | | | | | |
| | 13 | Reserved (was Concealment Enable) | | | | | | | | | |
| | 12 | Reserved (was Concealment Reference) | | | | | | | | | |
| | 11 | Reserved (was Concealment Type) | | | | | | | | | |
| | 10:9 | <p>Picture Coding Type. This field identifies whether the picture is an intra-coded picture (I), predictive-coded picture (P) or bi-directionally predictive-coded picture (B). See <i>ISO/IEC 13818-2</i> §6.3.9 for details.</p> <p>Format = MPEG_PICTURE_CODING_TYPE 00 = Reserved 01 = MPEG_I_PICTURE 10 = MPEG_P_PICTURE 11 = MPEG_B_PICTURE</p> | | | | | | | | | |
| | 8:1 | Reserved (was Slice Error Control) | | | | | | | | | |
| 0 | Reserved. (was Disable Mismatch) | | | | | | | | | | |



| Dword | Bit | Description |
|-------|------|---|
| 2 | 31:0 | <p>Interface Descriptor Remap Table [7:0]. This field contains the interface descriptor remap table entries for the first 8 kernel indices. Each table entry has 4 bits, providing a remapping range of [0, 15].</p> <p>This field is applicable to both frame picture destination (Picture Structure = 11) and field picture destination (Picture Structure = 01 or 10).</p> <p>Bits 31:28: Remap for index = 7 Bits 27:24: Remap for index = 6 Bits 23:20: Remap for index = 5 Bits 19:16: Remap for index = 4 Bits 15:12: Remap for index = 3 Bits 11:8: Remap for index = 2 Bits 7:4: Remap for index = 1 Bits 3:0: Remap for index = 0</p> <p>[DevCL] Errata: This field is reserved.</p> |
| 3 | 31:0 | <p>Interface Descriptor Remap Table [15:8]. This field contains the interface descriptor remap table entries for the last 8 kernel indices. Each table entry has 4 bits, providing a remapping range of [0, 15].</p> <p>This field is only applicable to frame destination. It is ignored when the destination is a field picture.</p> <p>[DevCL] Errata: This field is reserved.</p> |



1.6.4 VFE_STATE_EX [DevCTG+]

| Dword | Bit | Description |
|-------|------|---|
| | 7:0 | Reserved : MBZ |
| 1 | 31:0 | <p>VFE Control. This field is used by VFE depending on the mode of operation. See the following tables for details.</p> <p>If VFE Mode = AVC-IT or AVC-MC, this field is valid as defined in Error! Reference source not found.</p> <p>If VFE Mode = VC1-IT, this field is valid as defined in Error! Reference source not found.</p> <p>Otherwise, this field is reserved.</p> |
| 2 | 31:0 | <p>Interface Descriptor Remap Table. This field contains the interface descriptor remap table entries for the first 8 kernel indices. Each table entry has 4 bits, providing a remapping range of [0, 15].</p> <p>The input of this table is the Interface Descriptor Offset within the MEDIA_OBJECT or MEDIA_OBJECT_EX command. As the table is limited to map the first 16 values, any Interface Descriptor Offset greater than 15 is not remapped.</p> <p>Bits 31:28: Remap for index = 7 Bits 27:24: Remap for index = 6 Bits 23:20: Remap for index = 5 Bits 19:16: Remap for index = 4 Bits 15:12: Remap for index = 3 Bits 11:8: Remap for index = 2 Bits 7:4: Remap for index = 1 Bits 3:0: Remap for index = 0</p> |
| 3 | 31:0 | <p>Interface Descriptor Remap Table (cont). This field contains the interface descriptor remap table entries for the next 8 kernel indices (index = 8...15). Each table entry has 4 bits, providing a remapping range of [0, 15].</p> <p>Bits 31:28: Remap for index = 15 Bits 27:24: Remap for index = 14 Bits 23:20: Remap for index = 13 Bits 19:16: Remap for index = 12 Bits 15:12: Remap for index = 11 Bits 11:8: Remap for index = 10 Bits 7:4: Remap for index = 9 Bits 3:0: Remap for index = 8</p> |
| 4 | 31 | <p>Scoreboard Enable. This field enables and disables the hardware scoreboard in the Media Pipeline. If this field is cleared, hardware ignores the following scoreboard state fields.</p> <p>0 – Scoreboard disabled 1 – Scoreboard enabled [DevCTG] Reserved : MBZ</p> |



| Dword | Bit | Description |
|-------|-------|---|
| | 30 | Scoreboard Type. This field selects the type of scoreboard in use. 0 – Stalling scoreboard 1 – Non-stalling scoreboard [DevCTG] Reserved : MBZ [DevILK] This field must be zero (stalling scoreboard) |
| | 29:8 | Reserved : MBZ |
| | 7:0 | Scoreboard Mask. Each bit indicates the corresponding dependency scoreboard is enabled. The scoreboard is based on the relative (X, Y) distance from the current threads' (X, Y) position. Bit n (for n = 0...7): Score n is enabled Format = TRUE/FALSE [DevCTG] Reserved : MBZ |
| 5 | 31:28 | Scoreboard 3 Delta Y. Relative vertical distance of the dependent instance assigned to scoreboard 3, in the form of 2's compliment. Format = S3 [DevCTG] Reserved : MBZ |
| | 27:24 | Scoreboard 3 Delta X. Relative horizontal distance of the dependent instance assigned to scoreboard 3, in the form of 2's compliment. Format = S3 [DevCTG] Reserved : MBZ |
| | 23:16 | Scoreboard 2 Delta (X, Y) [DevCTG] Reserved : MBZ |
| | 15:8 | Scoreboard 1 Delta (X, Y) [DevCTG] Reserved : MBZ |
| | 7:0 | Scoreboard 0 Delta (X, Y) [DevCTG] Reserved : MBZ |
| 6 | 31:24 | Scoreboard 7 Delta (X, Y) [DevCTG] Reserved : MBZ |
| | 23:16 | Scoreboard 6 Delta (X, Y) [DevCTG] Reserved : MBZ |
| | 15:8 | Scoreboard 5 Delta (X, Y) [DevCTG] Reserved : MBZ |
| | 7:0 | Scoreboard 4 Delta (X, Y) [DevCTG] Reserved : MBZ |
| 7 | 31:0 | Reserved : MBZ |



VFE Control in AVC-IT or AVC-MC Mode

| Bit | Description |
|-------|--|
| 31:27 | Reserved : MBZ |
| 26 | <p>Residual Data Fixed Source Offset Flag</p> <p>0: Residual data are packed right behind the Motion Vector and Weight-Offset blocks</p> <p>1: Residual data start at the fixed offset provided by Residual Data Source Offset field.</p> <p>This field must be 1 if VFE_Mode = AVC-MC</p> |
| 25:24 | <p>Indirect Sub-Field Present Flag. This field indicates if any sub-field in the indirect data buffer is present. NoMV may only be used for an I-picture, where all macroblocks must have MvSize = 0.</p> <p>00: NoMV. Motion Vector and Weight/Offset fields are not present in the indirect data buffer.</p> <p>01: NoWO. Motion Vectors may be present in the indirect data buffer, but Weight/Offset is not.</p> <p>10: Reserved</p> <p>11: Both Motion Vector field and Weight/Offset field may be present</p> |
| 23:16 | <p>Residual Data Source Offset. This field specifies the fixed distance of the residual data from Indirect Data Start Address.</p> <p>It is in unit of dwords and must be 8-dword aligned.</p> <p>It is only valid when Residual Data Fixed Source Offset Flag is set and must be zero otherwise.</p> |
| 15:13 | Reserved : MBZ |
| 12:8 | <p>Inter Weight-Offset Residual Data GRF Destination Offset: This field specifies the offset in unit of GRF registers of the Weight/Offset data relative to the leading GRF location where the indirect data (where Motion Vectors, if present, are stored) will be placed. This field doesn't apply to intra-predicted (or I_PCM) macroblocks. In other words, this field is ignored by hardware for macroblocks with MvSize = 0.</p> <p>This field must be programmed to be greater than or equal to the maximum size for Motion Vectors data block. This field must be zero if Indirect Sub-Field Present Flag is NoMV.</p> |
| 7:5 | Reserved : MBZ |
| 4:0 | <p>Inter Residual Data GRF Destination Offset: This field specifies the offset in unit of GRF registers of the residual data of an inter-predicted macroblock relative to the leading GRF location of the indirect data (where Motion Vectors, if present, are stored).</p> <p>This field must be programmed to be greater than or equal to the maximum size for Motion Vectors and Weight-Offset data blocks. Therefore, it must be zero if Indirect Sub-Field Present Flag is NoMV.</p> <p>Note that this field doesn't affect intra-predicted macroblocks (or I_PCM macroblocks). As there is no motion vector for an I-macroblock (or I_PCM), the residual data will be placed at the beginning of the indirect data GRF location. In other words, this field is ignored by hardware for macroblocks with MvSize = 0.</p> |



VFE Control in VC1-IT Mode

| Bit | Description |
|------|----------------|
| 31:0 | Reserved : MBZ |

1.6.4.1 Interface Descriptor Remapping

When Extended VFE State VFE_STATE_EX is selected, the Interface Descriptor Offset field of the primitive command goes through the Interface Descriptor Remap Table, when Interface Descriptor Offset is within [0, 15]. When no Extended VFE State is selected, the interface descriptor remap is bypassed.

1.6.5 INTERFACE_DESCRIPTOR

| DWord | Bit | Description |
|-------|-------|--|
| 0 | 31:6 | Kernel Start Pointer. Specifies the 64-byte aligned address offset of the first instruction in the kernel. This pointer is relative to the General State Base Address [Pre-DevILK] or Instruction Base Address [Dev ILK] [DevBW] Errata BWT007: Instructions pointed at by offsets from General State Base must be contained within 32-bit physical address space (that is, must map to memory pages under 4G.) [Pre-DevILK]: Format = GeneralStateOffset[31:6] [DevILK+]: Format = InstructionBaseOffset[31:6] |
| | 5:4 | Reserved : MBZ |
| | 3:0 | GRF Register Blocks [Pre-DevSNB]. Defines the number of GRF Register Blocks used by the kernel. A register block contains 8 registers. A kernel using a register count that is not a multiple of 8 must round up to the next multiple of 8. Format = U4 register block count - 1 Range = [0,15] corresponding to [1,16] 8-register blocks Restriction: LSB must be zero, indicating that GRF assignment is in granularity of 16 GRF registers. |
| 1 | 31:26 | Constant URB Entry Read Length. Specifies the amount of URB data read and passed in the thread payload for the Constant URB entry, in 8-DW register increments. A value 0 means that no Constant URB Entry will be loaded. The Constant URB Entry Read Offset field will then be ignored. Format = U6 Range = [0,63] |
| | 25:20 | Constant URB Entry Read Offset. Specifies the offset (in 8-DW units) at which Constant URB data is to be read from the URB before being included in the thread payload. Format = U6 Range = [0,63] |
| | 19 | Reserved : MBZ |
| | 18 | Single Program Flow (SPF). Specifies whether the kernel program has a single program flow (SIMDn _{xm} with m = 1) or multiple program flows (SIMDn _{xm} with m > 1). 0: Multiple Program Flows 1: Single Program Flow |



| DWord | Bit | Description |
|-------|-------|---|
| | 17 | <p>Thread Priority. Specifies the priority of the thread for dispatch</p> <p>0: Normal Priority 1: High Priority</p> <p>Programming Notes:</p> <ul style="list-style-type: none"> [Pre-DevILK]: This field must be set to zero. |
| | 16 | <p>Floating Point Mode. Specifies the floating point mode used by the dispatched thread.</p> <p>0: Use IEEE-754 Rules 1: Use alternate rules</p> |
| | 15:14 | Reserved: MBZ |
| | 13 | <p>Illegal Opcode Exception Enable. This bit gets loaded into EU CR0.1[12] (note the bit # difference). See <i>Exceptions</i> and <i>ISA Execution Environment</i>.</p> <p>Format: Enable</p> |
| | 12 | Reserved: MBZ |
| | 11 | <p>MaskStack Exception Enable. This bit gets loaded into EU CR0.1[11]. See <i>Exceptions</i> and <i>ISA Execution Environment</i>.</p> <p>Format: Enable</p> |
| | 10:8 | Reserved: MBZ |
| | 7 | <p>Software Exception Enable. This bit gets loaded into EU CR0.1[13] (note the bit # difference). See <i>Exceptions</i> and <i>ISA Execution Environment</i>.</p> <p>Format: Enable</p> |
| | 6:0 | Reserved: MBZ |
| 2 | 31:5 | <p>Sampler State Pointer. Specifies the 32-byte aligned address offset of the sampler state table. This pointer is relative to the General State Base Address.</p> <p>[DevBW-A] Errata BWT007: Sampler state pointed at by offsets from General State Base must be contained within 32-bit physical address space (that is, must map to memory pages under 4G.)</p> <p>Format = GeneralStateOffset[31:5] <i>This field is ignored for child threads.</i></p> |
| | 4:2 | <p>Sampler Count. Specifies how many samplers (in multiples of 4) the kernel uses. Used only for prefetching the associated sampler state entries.</p> <p>Format = U3 Range = [0,4]</p> <p>0: no samplers used 1: between 1 and 4 samplers used 2: between 5 and 8 samplers used 3: between 9 and 12 samplers used 4: between 13 and 16 samplers used</p> <p><i>This field is ignored for child threads.</i> <i>If this field is not zero, sampler state is prefetched for the first instance of a root thread upon the startup of the media pipeline.</i></p> |
| | 1:0 | Reserved : MBZ |



| DWord | Bit | Description |
|-------|------|--|
| 3 | 31:5 | Binding Table Pointer. Specifies the 32-byte aligned address of the binding table. This pointer is relative to the Surface State Base Address . Format = SurfaceStateOffset[31:5] <i>This field is ignored for child threads.</i> |
| | 4:0 | Binding Table Entry Count. Specifies how many binding table entries the kernel uses. Used only for prefetching of the binding table entries and associated surface state. Note: The maximum number of prefetched binding table entries is limited to 31. For kernels using a large number of binding table entries, it may be wise to set this field to zero to avoid prefetching too many entries and thrashing the state cache. Format = U5 Range = [0,31] <i>This field is ignored for child threads.</i> <i>If this field is not zero, binding table and surface state are prefetched for the first instance of a root thread upon the startup of the media pipeline.</i> |



1.7 Media State and Primitive Commands

1.7.1 MEDIA_STATE_POINTERS Command ([Pre-DevSNB])

The MEDIA_STATE_POINTERS command is used to set up the pointers to the VFE states (VFE state, VLD state or VFE state extension). This command is issued prior to a set of media primitive commands, and points to the Generic mode VFE state and VLD decode mode VLD state (or VFE extended state).

[DevBW] Errata BWT007: State data pointed at by offsets from General State Base must be contained within 32-bit physical address space (that is, must map to memory pages under 4G.)

| DWord | Bit | Description |
|-------|-------|--|
| 0 | 31:29 | Command Type = GFXPIPE = 3h |
| | 28:16 | Media Command Opcode = MEDIA_STATE_POINTERS Pipeline[28:27] = Media = 2h; Opcode[26:24] = 0h; Subopcode[23:16] = 0h |
| | 15:0 | DWord Length (Excludes DWords 0,1) = 01h |
| 1 | 31:5 | Extended State Pointer. Specifies the 32-byte aligned address of the extended VFE state (either VLD_STATE or VFE_STATE_EX). This pointer is relative to the General State Base Address . Which extended VFE state is used depends on VFE Mode . If VFE Mode is set to VLD Mode (0001), VLD_STATE is used. Otherwise, VFE_STATE_EX is used. Format = GeneralStateOffset[31:5] [DevBW, DevCL] Note that VFE_STATE_EX is reserved |
| | 4:1 | Reserved : MBZ |
| | 0 | Extended VFE State Enable (was VLD Enable). This field specifies whether extended VFE state is loaded. 0 = Disabled. No extended VFE state (and Extension State Pointer is ignored). 1 = Enabled. The extended VFE state pointed by Extended State Pointer is loaded [DevBW, DevCL] Note that VFE_STATE_EX is reserved |
| 2 | 31:5 | Pointer to VFE_STATE. Specifies the 32-byte aligned address of the VFE_STATE. This pointer is relative to the General State Base Address . Format = GeneralStateOffset[31:5] |
| | 4:0 | Reserved : MBZ |



1.7.2 MEDIA_OBJECT Command ([Pre-DevGT])

The MEDIA_OBJECT command is the basic media primitive command for the media pipeline. It supports loading of inline data as well as indirect data.

The MEDIA_OBJECT command can be used in the following three VFE modes: Generic mode, IS mode and VLD mode.

The MEDIA_OBJECT command cannot be used in the following VFE modes: AVC-IT, AVC-MC, and VC1-IT.

| Dword | Bits | Description |
|-------|-------|--|
| 0 | 31:29 | Command Type = GFXPIPE = 3h |
| | 28:16 | Media Command Opcode = MEDIA_OBJECT Pipeline[28:27] = Media = 2h; Opcode[26:24] = 1h; Subopcode[23:16] = 0h |
| | 15:0 | DWord Length (Excludes DWords 0,1) VLD Mode: DWord Length = 4. There are 2 DW of inline data in this mode. IS Mode: DWord Length = N+2, where N is the number of DW of inline data (N >= 10). According to the inline format table shown in the following section, N is 10. However, hardware must be able to handle different size of N, as software may determine later to transfer additional driver/kernel information inline. Generic Mode: DWord Length = N+2, where N is in the range of [0,504]. The maximum is 504 DW (equivalent to 63 8-DW registers). When both inline and indirect data are fetched for this command, the total size in 8-DW registers must be less than or equal to 63 (with both inline data length N and indirect data length rounded up to 8-DW aligned individually). If indirect data are fetched, the minimal inline data length is 0. If indirect data are not fetched, the minimal inline data is 1DW. Note: Regardless of the mode, inline data must be present in this command. |
| | 7 | Reserved. MBZ |
| | 6:0 | Interface Descriptor Offset. This field specifies the offset from the interface descriptor base pointer to the interface descriptor which will be applied to this object. It is specified in units of interface descriptors. <i>In VLD mode, this field is ignored by hardware.</i> Format = U7 |
| 2 | 31:29 | Reserved. MBZ |
| | 28 | Retain Bit. The hardware will keep the last 256-bit quantity of this indirect object in-use after the object transfer is complete. A subsequent Indirect object packet will use the retained 256bit quantity as the first piece of data. Format = Enable (1) /Disable (0) [DevBW] Erratum: this field is reserved: MBZ |
| | 27:25 | Reserved. MBZ |
| | 24 | Thread Synchronization. This field when set indicates that the dispatch of the thread originated from this command is based on the "spawn root thread" message. <i>In VLD mode, this field must be programmed as 0, because the Children Present field in VFE_STATE must be 0 in this mode.</i> 0 = No thread synchronization 1 = Thread dispatch is synchronized by the "spawn root thread" message |



| Dword | Bits | Description |
|-------|-------|--|
| | 23:17 | Reserved. MBZ |
| | 16:0 | <p>Indirect Data Length. This field provides the length in bytes of the indirect data. A value zero indicates that indirect data fetching is disabled – subsequently, the Indirect Data Start Address field is ignored.</p> <p>This field must have the same alignment as the Indirect Object Data Start Address.</p> <p>VLD Mode: It is the length in bytes of the bitstream data for the current slice. It includes the first byte of the first macroblock and the last non-zero byte of the last macroblock in the slice. Specifically, the zero-padding bytes (if present) and the next start-code are excluded. Hardware ignores the contents after the last non-zero byte. This field is sized to support MPEG-2 MP@HL bitstream. According to Table 8-6 of <i>ISO/IEC 13818-2</i>, the maximum number of bits per macroblock for 4:2:0 is 4608. So the maximum slice size for MP@HL (e.g. 1080i) is $4608 * 120 / 8 = 69120$ bytes (0x10E00), which requires 17 bits. Programming Restriction: hardware has the limitation that the maximum allowed value is 0x1FFE0. As MPEG-2 spec does not post any limitation of the size of zero-padding bytes, it is possible to have a slice data with large length (including zero-padding bytes). As the data beyond 0x10E00 would only be zero bytes for a valid slice data, it is recommended that host software truncates the indirect data length to, say, 1x12000. Hardware can take care of the zero-padding bytes beyond the last non-zero byte of the slice.</p> <p>IS Mode: It must be DWord aligned.</p> <p>Generic Mode: It must be DQWord (32-byte) aligned. As the indirect data are sent directly to URB, range is limited to 496 DW. When both inline and indirect data are fetched for this command, the total size in 8-DW registers must be less than or equal to 63 (with both inline data length and indirect data length rounded up to 8-DW aligned).</p> <p>[DevBW-A] Erratum: In Generic Mode, the length alignment restrict is relaxed to be DWord alignment..</p> <p>Format = U17 in bytes</p> |
| 3 | 31:0 | <p>Indirect Data Start Address. This field specifies the Graphics Memory starting address of the data to be loaded into the kernel for processing. This pointer is relative to the Indirect Object Base Address.</p> <p>Hardware ignores this field if indirect data is not present.</p> <p>Alignment of this address depends on the mode of operation.</p> <p>VLD Mode: It is the byte aligned address for the VLD bitstream data.</p> <p>IS Mode: It is the DWord aligned address for the first IDCT coefficients.</p> <p>Generic Mode: It is the DWord aligned address of the indirect data.</p> <p>Range = [0 - 512MB] (Bits 31:29 MBZ)</p> |
| 4..N | 31:0 | <p>Inline Data</p> <p>IS and VLD Modes: Hardware interprets this data in the specified format.</p> <p>Generic Mode: The format of this data is specified by software. Hardware does not interpret this data; it merely passes it to the kernel for processing. The total size for the inline data and indirect data must not exceed the URB allocation size.</p> |



1.7.2.1 Inline and Indirect Data Format in Generic Mode

In Generic mode, inline data must be present. All inline data will be delivered to the thread's payload starting and ending on the 8-DW aligned register boundary. Inline data starts on dword 4 of the MEDIA_OBJECT command. If the dword length field of the MEDIA_OBJECT command is N+2, the size of the inline data will N. VFE always zero-pads inline data into 8-DW before delivering to URB. If N is multiple of 8-DW, the inline data corresponds to exactly N/8 GRF registers. If N is not multiple of 8-DW, there will be (N/8 + 1) registers written for the inline data with the last register containing the last a few dwords of inline data with remaining dwords zeroed out by VFE.

Indirect data, if present, will be written into GRF registers immediately following the inline data in the thread's payload. Alignment and padding for indirect data are the same as that for inline data. In short, indirect data are also starting and ending on 8-DW aligned register boundary. If indirect data length is not multiple of 8-DW, VFE hardware will zero pad the last GRF register.

1.7.2.2 Inline and Indirect Data Format in IS Mode

Each MEDIA_OBJECT command in "IS mode" corresponds to the processing of one macroblock. Macroblock parameters are passed in as inline data and the non-zero DCT coefficient data for the macroblock is passed in as indirect data.

Error! Reference source not found. depicts the inline data format in IS mode. All fields in inline data are forwarded to the thread as thread payload. Alignment and padding is identical to that described for Generic mode. Some fields are merely forwarded. Some fields are also used by VFE as indicated in the following table by a mark of [Used by VFE]. As shown, inline data starts at dword 4 of MEDIA_OBJECT command. There are 10 dwords total.

Inline data in IS mode

| DWord | Bit | Description | | | | | | | | | | | | | | | | | | | | |
|-------|---|---|-------------|-------------|---------------------------------|---------------------------------|----|---|---|---|----|---|---|---|----|---|---|---|----|---|---|---|
| 4+0 | 31:28 | <p>Motion Vertical Field Select. A bit-wise representation of a long [2][2] array as defined in §6.3.17.2 of the <i>ISO/IEC 13818-2</i> (see also §7.6.4).</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>MVector [r]</th> <th>MVector [s]</th> <th>MotionVerticalFieldSelect Index</th> </tr> </thead> <tbody> <tr> <td>28</td> <td>0</td> <td>0</td> <td>0</td> </tr> <tr> <td>29</td> <td>0</td> <td>1</td> <td>1</td> </tr> <tr> <td>30</td> <td>1</td> <td>0</td> <td>2</td> </tr> <tr> <td>31</td> <td>1</td> <td>1</td> <td>3</td> </tr> </tbody> </table> <p>Format = MC_MotionVerticalFieldSelect. 0 = The prediction is taken from the <u>top</u> reference field. 1 = The prediction is taken from the <u>bottom</u> reference field.</p> | Bit | MVector [r] | MVector [s] | MotionVerticalFieldSelect Index | 28 | 0 | 0 | 0 | 29 | 0 | 1 | 1 | 30 | 1 | 0 | 2 | 31 | 1 | 1 | 3 |
| | | Bit | MVector [r] | MVector [s] | MotionVerticalFieldSelect Index | | | | | | | | | | | | | | | | | |
| | | 28 | 0 | 0 | 0 | | | | | | | | | | | | | | | | | |
| 29 | 0 | 1 | 1 | | | | | | | | | | | | | | | | | | | |
| 30 | 1 | 0 | 2 | | | | | | | | | | | | | | | | | | | |
| 31 | 1 | 1 | 3 | | | | | | | | | | | | | | | | | | | |
| 27 | <p>Second Field. This bit indicates that this is the second field in the current frame. The prediction for this macroblock, if it belongs to a field P-picture, should use this bit to determine which frame contains the reference field as described in §7.6.2.1 of the <i>ISO/IEC 13818-2</i>.</p> <p>When the picture type is not P or the prediction type is not field, this value should be 0.</p> <p>Format = MC_SecondPField 0 = This is not the second field. 1 = This is the second field.</p> | | | | | | | | | | | | | | | | | | | | | |
| 26 | <p>Reserved. (HWMC mode)</p> | | | | | | | | | | | | | | | | | | | | | |



| DWord | Bit | Description | | | | | | | | | | | | | | | |
|-------|---|---|-------|---|--|------|----------|----------|------|-------|-------|------|-------|------|------|------------|------------|
| | 25:24 | <p>Motion Type. When combined with the destination picture type (field or frame) this Motion Type field indicates the type of motion to be applied to the macroblock. See <i>ISO/IEC 13818-2</i> §6.3.17.1, Tables 6-17, 6-18. In particular, the device supports dual-prime motion prediction (11) in both frame and field picture type.</p> <p>Format = MC_MotionType</p> <table border="1" data-bbox="623 506 1357 711"> <thead> <tr> <th>Value</th> <th>Destination = Frame Picture_Structure = 11</th> <th>Destination = Field Picture_Structure != 11</th> </tr> </thead> <tbody> <tr> <td>'00'</td> <td>Reserved</td> <td>Reserved</td> </tr> <tr> <td>'01'</td> <td>Field</td> <td>Field</td> </tr> <tr> <td>'10'</td> <td>Frame</td> <td>16x8</td> </tr> <tr> <td>'11'</td> <td>Dual-Prime</td> <td>Dual-Prime</td> </tr> </tbody> </table> | Value | Destination = Frame Picture_Structure = 11 | Destination = Field Picture_Structure != 11 | '00' | Reserved | Reserved | '01' | Field | Field | '10' | Frame | 16x8 | '11' | Dual-Prime | Dual-Prime |
| Value | Destination = Frame Picture_Structure = 11 | Destination = Field Picture_Structure != 11 | | | | | | | | | | | | | | | |
| '00' | Reserved | Reserved | | | | | | | | | | | | | | | |
| '01' | Field | Field | | | | | | | | | | | | | | | |
| '10' | Frame | 16x8 | | | | | | | | | | | | | | | |
| '11' | Dual-Prime | Dual-Prime | | | | | | | | | | | | | | | |
| | 23:22 | Reserved. (Scan method) | | | | | | | | | | | | | | | |
| | 21 | <p>DCT Type. This field specifies the DCT type of the current macroblock. The kernel should ignore this field when processing Cb/Cr data. See <i>ISO/IEC 13818-2</i> §6.3.17.1. This field is zero if Coded Block Pattern is also zero (no coded blocks present).</p> <p>0 = MC_FRAME_DCT (Macroblock is frame DCT coded). 1 = MC_FIELD_DCT (Macroblock is field DCT coded).</p> | | | | | | | | | | | | | | | |
| | 20 | Overlap Transform (H261 Loop Filter). This field, when set, indicates that overlap smoothing filter is performed after motion compensation and before in-loop deblocking. | | | | | | | | | | | | | | | |
| | 19 | <p>4MV Mode: (H263/WMV)</p> <p>This field indicates if the current macroblock is coded with 4 motion vectors, one for each 8x8 block.</p> | | | | | | | | | | | | | | | |
| | 18 | <p>Macroblock Motion Backward. This field specifies if the backward motion vector is active. See <i>ISO/IEC 13818-2</i> Tables B-2 through B-4.</p> <p>0 = No backward motion vector. 1 = Use backward motion vector(s).</p> | | | | | | | | | | | | | | | |
| | 17 | <p>Macroblock Motion Forward. This field specifies if the forward motion vector is active. See <i>ISO/IEC 13818-2</i> Tables B-2 through B-4.</p> <p>0 = No forward motion vector. 1 = Use forward motion vector(s).</p> | | | | | | | | | | | | | | | |
| | 16 | <p>Macroblock Intra Type. This field specifies if the current macroblock is intra-coded. When set, Coded Block Pattern is ignored and no prediction is performed (i.e., no motion vectors are used). See <i>ISO/IEC 13818-2</i> Tables B-2 through B-4.</p> <p>0 = Non-intra macroblock. 1 = Intra macroblock.</p> | | | | | | | | | | | | | | | |
| | 15:0 | Reserved. (MB address) | | | | | | | | | | | | | | | |
| 4+1 | 31:24 | Reserved. (Skip Macroblocks) | | | | | | | | | | | | | | | |
| | 23:0 | Reserved. (Offset into error data) | | | | | | | | | | | | | | | |
| 4+2 | 31:24 | Subblock Coding for Block Y1 | | | | | | | | | | | | | | | |



| DWord | Bit | Description |
|-------|-------|--|
| | 23:16 | Subblock Coding for Block Y0. This field specifies the subblock partition and subblock coding pattern for the block. The definition of the 8 bits of this field is listed below. Detailed coding can be found in Error! Reference source not found.. Bits [7:6]: reserved Bits [5:2]: Subblock present Bits [1:0]: Subblock partitioning |
| | 15:12 | Reserved. |
| | 11:6 | Coded Block Pattern. This field specifies whether blocks are present or not. Format = 6-bit mask. Bit 11: Y0 Bit 10: Y1 Bit 9: Y2 Bit 8: Y3 Bit 7: Cb4 Bit 6: Cr5 [Used by VFE] |
| | 5:0 | Reserved. (Quantization Scale Code) |
| 4+3 | 31:24 | Subblock Coding for Block Cr5 |
| | 23:16 | Subblock Coding for Block Cb4 |
| | 15:8 | Subblock Coding for Block Y3 |
| | 7:0 | Subblock Coding for Block Y2 |
| 4+4 | 31:16 | Motion Vectors – Field 0, Forward, Vertical Component. Each vector component is a 16-bit two's-complement value. The vector is relative to the current macroblock location. According to ISO/IEC 13818-2 Table 7-8, the valid range of each vector component is [-2048, +2047.5], implying a format of s11.1. However, it should be noted that motion vector values are sign extended to 16 bits. |
| | 15:0 | Motion Vectors – Field 0, Forward, Horizontal Component |
| 4+5 | 31:16 | Motion Vectors – Field 0, Backward, Vertical Component |
| | 15:0 | Motion Vectors – Field 0, Backward, Horizontal Component |
| 4+6 | 31:16 | Motion Vectors – Field 1, Forward, Vertical Component |
| | 15:0 | Motion Vectors – Field 1, Forward, Horizontal Component |
| 4+7 | 31:16 | Motion Vectors – Field 1, Backward, Vertical Component |
| | 15:0 | Motion Vectors – Field 1, Backward, Horizontal Component |
| 4+8 | 31:30 | Reserved. |
| | 29 | May need this for WMV. (Interpolation Rounder Control) |
| | 28 | May need this for WMV. (Bidirectional Averaging Control) |
| | 27:20 | Reserved. |



| DWord | Bit | Description |
|-------|-------|---|
| | 19:18 | Picture Coding Type. This field identifies whether the picture is an intra-coded picture (I), predictive-coded picture (P) or bi-directionally predictive-coded picture (B). See <i>ISO/IEC 13818-2</i> §6.3.9 for details. Format = MPEG_PICTURE_CODING_TYPE 00 = Reserved 01 = MPEG_I_PICTURE 10 = MPEG_P_PICTURE 11 = MPEG_B_PICTURE |
| | 17:16 | Picture Structure. This field specifies whether the picture is encoded in the form of a frame picture or one field (top or bottom) picture. See <i>ISO/IEC 13818-2</i> §6.3.10 for details. Format = MPEG_PICTURE_STRUCTURE 00 = Reserved 01 = MPEG_TOP_FIELD 10 = MPEG_BOTTOM_FIELD 11 = MPEG_FRAME |
| | 15 | Reserved. (8-bit Intra) |
| | 14:13 | Reserved. (Intra DC Precision) |
| | 12:0 | Reserved. |
| 4+9 | 31:27 | Reserved. |
| | 26:20 | Vertical Origin. Set the vertical origin of the next macroblock in the destination picture in units of macroblocks. (Valid range is 0 to 120). Format = U7 in macroblock units. Range = [0, 120] |
| | 19:11 | Reserved: MBZ |
| | 10:4 | Horizontal Origin. Set the horizontal origin of the next macroblock in the destination picture in units of macroblocks. Format = U7 in macroblock units. Range = [0, 127] |
| | 3:0 | Reserved. |

The control parameters for inverse-scan are carried in the inline data packet. In particular, the Coded Block Pattern field in DW6 is used to determine how many blocks are coded and therefore how many blocks will be output from the inverse-scan.

Besides that dword 6 (containing Coded Block Pattern field) is used by VFE hardware as control parameter for inverse-scan, the rest of the inline data are determined between the host software and the kernel software. Therefore, the exact format and size of the inline data may differ, as long as the Coded Block Pattern field in dword 6 remains the same.

Error! Reference source not found. shows a ‘recommended’ inline data format for WMV9_C support. Support for other video coding standard may take a similar format. For example, bits [16:31] of DW6 as well as DW7 carry the subblock coding information for WMV9 acceleration. According to WMV video specification, each 8x8 block may be further subdivided into two 8x4, two 4x8, or four 4x4 subblocks. For each block, the subblock coding occupies one byte with only the lower 6 bits utilized according to **Error! Reference source not found.** For MPEG-2 IDCT support, as there is no block subdivision, these fields may be reserved, or used to carry other information.

Subblock coding (bits [7:6] are reserved).

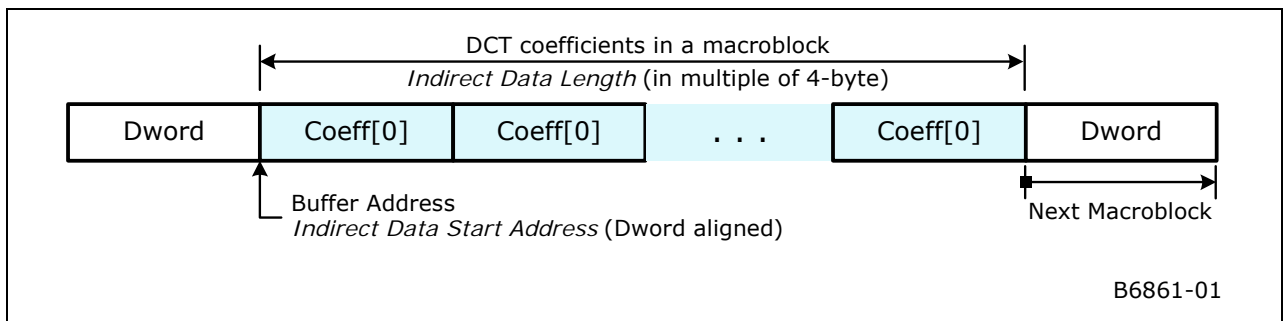
| Subblock Partitioning (Bits [1:0]) | | Subblock Present (0 means not present, 1 means present) | | | |
|------------------------------------|----------------------------|---|------------|------------|------------|
| Code | Meaning | Bit 2 | Bit 3 | Bit 4 | Bit 5 |
| 00 | Single 8x8 block (sb0) | Sb0 | Don't care | Don't care | Don't care |
| 01 | Two 8x4 subblocks (sb0-1) | Sb0 | Sb1 | Don't care | Don't care |
| 10 | Two 4x8 subblocks (sb0-1) | Sb0 | Sb1 | Don't care | Don't care |
| 11 | Four 4x4 subblocks (sb0-3) | Sb0 | Sb1 | Sb2 | Sb3 |

The block data output from the inverse scan will follow immediately after the inline data in the thread's payload, again, aligning to GRF register. Block data output by nature is 8-DW aligned. The actual size depends on the coded block pattern. As each block contains 8x8 16-bit DCT coefficients, if the total number of coded block is M, the block data will take $8 * 8 * 2 * M / 32 = 4 * M$ GRE registers.

As VFE performs inverse-scan on the indirect data, the indirect data must follow the exact format described in Figure 1-4 and

The indirect data start address in MEDIA_OBJECT specifies the doubleword aligned address of the first non-zero DCT coefficient of the first block of the macroblock. Only the non-zero DCT coefficients are present in the data buffer and they are packed in the block sequence of Y0, Y1, Y2, Y3, Cb4 and Cr5, as shown in Figure 1-4. The indirect data length in MEDIA_OBJECT includes all the non-zero coefficients for the macroblock. It must be doubleword aligned.

Figure 1-4. Structure of the IDCT Compressed Data Buffer



Each non-zero coefficient in the indirect data buffer is contained in a doubleword-size data structure containing the coefficient index, end of block (EOB) flag and the fixed-point coefficient value in 2's complement form. As shown in

, *index* is the row major 'raster' index of the coefficient within an 8x8 block. DCT coefficient is a 16-bit value in 2's complement, which is clamped to a 12-bit signed value by the host. Effectively, bit 27 is the sign bit. However, as the kernel software consumes these data as 16-bit quantities any way, VFE simply forwards these exact 16-bit DCT coefficients to the thread's payload.



Structure of a DCT coefficient unit

| DWord | Bit | Description |
|-------|-------|--|
| 0 | 31:16 | DCT Coefficient Value. This field contains the value of the non-zero DCT coefficient in 2's compliment. |
| | 15:7 | Reserved: MBZ |
| | 6:1 | Index. This field specifies the raster-scan address (raw address) of the DCT coefficient within the 8x8 block. For example, coefficient at location (row, column) = (0, 0) has an index of 0; that at (2, 3) has an index of $2*8 + 3 = 19$. Format = U6 Range = [0, 63] |
| | 0 | EOB (End of Block). This field indicates whether the DCT coefficient is the last one of the current block. |

1.7.2.3 Inline and Indirect Data Format in VLD Mode

A MEDIA_OBJECT command in “VLD mode” is used to process a slice using the VFE hardware. Slice header parameters are passed in as inline data and the bitstream data for the slice is passed in as indirect data. Of the inline data, slice_horizontal_position and slice_vertical_position determines the location within the destination picture of the first macroblock in the slice.

| DWord | Bits | Description |
|-------|-------|--|
| 4 | 31 | Reserved. MBZ |
| | 30:24 | Slice Horizontal Position. This 7-bit field indicates the horizontal position (in macroblock units) of the first macroblock in the slice. Format = U7 in macroblocks |
| | 23 | Reserved. MBZ |
| | 22:16 | Slice Vertical Position. This 7-bit field indicates the vertical position (in macroblock units) of the first macroblock in the slice. Format = U7 in macroblocks |
| | 15 | Reserved. MBZ |
| | 14:8 | Macroblock Count. This 7-bit field indicates the number of macroblocks in the slice, including skipped macroblocks. |
| | 7:3 | Reserved. MBZ. |
| | 2:0 | First Macroblock Bit Offset. This field provides the bit offset of the first macroblock in the first byte of the input bitstream. Format = U3 |
| 5 | 31:29 | Reserved. MBZ. |

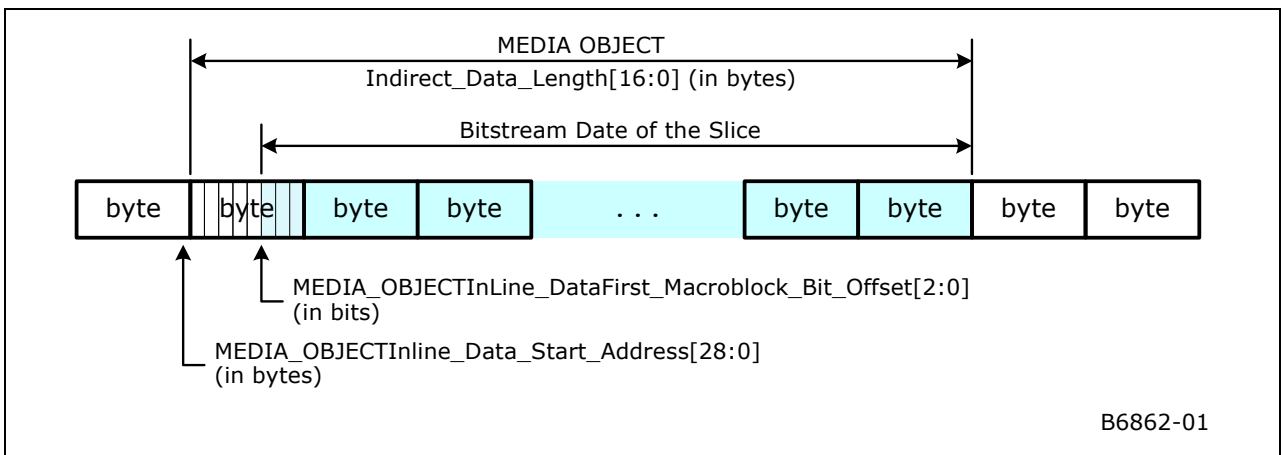


| DWord | Bits | Description |
|-------|-------|--|
| | 28:24 | Quantizer Scale Code. This field sets the quantizer scale code of the inverse quantizer. It remains in effect until changed by a decoded quantizer scale code in a macroblock. This field is decoded from the slice header by host software. Format = U5 (0 is Reserved) |
| | 23:0 | Reserved. MBZ. |

The indirect data start address in MEDIA_OBJECT specifies the starting Graphics Memory address of the bitstream data that follows the slice header. It provides the byte address for the first macroblock of the slice. Together with the First Macroblock Bit Offset field in the inline data, it provides the bit location of the macroblock within the compressed bitstream.

The indirect data length in MEDIA_OBJECT provides the length in bytes of the bitstream data for this slice. It includes the first byte of the first macroblock and the last **non-zero** byte of the last macroblock in the slice. Specifically, the zero-padding bytes (if present) and the next start-code are excluded. Hardware ignores the contents after the last non-zero byte. **Error! Reference source not found.** illustrates these parameters for a slice data.

Indirect data buffer for a slice



1.7.3 MEDIA_OBJECT_EX Command [DevCTG-DevILK]

[DevBW/DevCL] This command is not supported.

[DevBW/DevCL] This command is not supported.

The MEDIA_OBJECT_EX command is the extended media primitive command for the media pipeline. The command can be used for AVC and VC1 decode in different modes. It supports loading of inline data as well as indirect data.

This command can be used in the following VFE modes: AVC-IT, AVC-MC, VC1-IT and Generic modes.

This command cannot be used in the following VFE modes: IS and VLD modes.



| Dword | Bits | Description |
|-------|-------|---|
| 0 | 31:29 | Command Type = GFXPIPE = 3h |
| | 28:16 | Media Command Opcode = MEDIA_OBJECT_EX Pipeline[28:27] = Media = 2h; Opcode[26:24] = 1h; Subopcode[23:16] = 01h |
| | 15:0 | DWord Length (Excludes DWords 0,1) Note: Regardless of the mode, inline data must be present in this command. |
| 1 | 31:0 | VFE Dword. This field contains the data specific to the mode of operation. This field is consumed by VFE, and in general is not forwarded to the thread. |
| 2 | 31 | Reserved. MBZ |
| | 30:24 | Interface Descriptor Offset. This field specifies the offset from the interface descriptor base pointer to the interface descriptor which will be applied to this object. It is specified in units of interface descriptors. <i>Actual Interface Descriptor is selected post the Interface descriptor Remapping.</i> Format = U7 |
| | 23:22 | Reserved. MBZ |
| | 21 | [DevILK] Use Scoreboard. This field specifies whether the thread associated with this command uses hardware scoreboard. Only when this field is set, the scoreboard control fields in the VFE Dword are valid. If this field is cleared, the thread associated with this command bypasses hardware scoreboard. 0 = Not using scoreboard 1 = Using scoreboard [DevCTG] Reserved. MBZ |
| | 20 | Thread Synchronization. This field indicates whether the thread is a Synchronized Root Thread (SRT). 0 = Not a SRT 1 = SRT |
| | 19:17 | Reserved. MBZ |
| | 16:0 | Indirect Data Length. This field provides the length in bytes of the indirect data. A value zero indicates that indirect data fetching is disabled – subsequently, the Indirect Data Start Address field is ignored. This field must have the same alignment as the Indirect Object Data Start Address. AVC-IT Mode: It must be DWord aligned. AVC-MC Mode: It must be 8-DWord aligned. VC1-IT Mode: It must be DWord aligned. Generic Mode: It must be DWord aligned. Format = U17 in bytes |



| Dword | Bits | Description |
|-------|------|--|
| 3 | 31:0 | <p>Indirect Data Start Address. This field specifies the Graphics Memory starting address of the data to be loaded into the kernel for processing. This pointer is relative to the Indirect Object Base Address.</p> <p>Hardware ignores this field if indirect data is not present.</p> <p>Alignment of this address depends on the mode of operation.</p> <p>AVC-IT Mode: It is the DWord aligned address for the first field, if available, in the order of Motion Vectors, Weights/Offsets, transform-domain residual data (IDCT coefficients).</p> <p>AVC-MC Mode: It is the DWord aligned address for the first field, if available, in the order of Motion Vectors, Weights/Offsets, pixel-domain residual data.</p> <p>VC1-IT Mode: It is the DWord aligned address for the first IDCT coefficients.</p> <p>Generic Mode: It is the DWord aligned address.</p> <p>Range = [0 - 512MB] (Bits 31:29 MBZ)</p> |
| 4..N | 31:0 | <p>Inline Data</p> <p>AVC-IT Modes: Hardware interprets this data in the specified format.</p> <p>AVC-MC Modes: Hardware interprets this data in the specified format.</p> <p>VC1-IT Modes: Hardware interprets this data in the specified format.</p> <p>Generic Modes: Hardware does not interpret this data</p> |

[DevILK+] Up to eight dependencies are provided in Generic mode, which are explicitly controlled by software programming using the VFE Dword field.

The VFE Dword [DevCTG]

| DWord | Bit | Description |
|-------|------|----------------------|
| 1 | 31:0 | Reserved: MBZ |

The VFE Dword [DevILK]

| DWord | Bit | Description |
|-------|-------|--|
| 1 | 31:28 | <p>Scoreboard Mask 4-7 : Each bit indicates the corresponding dependency scoreboard is dependent on. This field is AND'd with the corresponding Scoreboard Mask field in the VFE_STATE_EX.</p> <p>Bit n (for n = 4...7): Scoreboard n is dependent, where bit 28 maps to n = 4.</p> <p>Format = TRUE/FALSE</p> |
| | 27:26 | <p>Scoreboard Color : This field specifies which dependency color the current thread belongs to. It affects the dependency scoreboard control.</p> <p>Format = U2</p> |
| | 25 | Reserved. MBZ |
| | 24:16 | <p>Scoreboard Y</p> <p>This field provides the Y term of the scoreboard value of the current thread.</p> <p>Format = U9</p> |



| DWord | Bit | Description |
|-------|-------|---|
| | 15:12 | <p>Scoreboard Mask 0-3 : Each bit indicates the corresponding dependency scoreboard is dependent on. This field is AND'd with the corresponding Scoreboard Mask field in the VFE_STATE_EX.</p> <p>Bit n (for n = 0...3): Scoreboard n is dependent, where bit 12 maps to n = 0. Format = TRUE/FALSE</p> |
| | 11:9 | Reserved. MBZ |
| | 8:0 | <p>Scoreboard X This field provides the X term of the scoreboard value of the current thread. Format = U9</p> |

1.7.3.1 Inline Data Format in AVC-IT Mode

Each MEDIA_OBJECT_EX command in “AVC-IT mode” corresponds to the processing of one macroblock. Macroblock parameters are passed in as inline data and the non-zero DCT coefficient data (as well as motion vectors and weight/offset) for the macroblock is passed in as indirect data.

Error! Reference source not found. depicts the inline data format in AVC-IT mode. All fields in inline data are forwarded to the thread as thread payload, except the QP fields, where the derived macroblock information is filled in. Starting at GRF location, inline data are stored in GRF contiguously with the tail-end partial GRF, if present, zero-filled. Some fields are merely forwarded. Some fields are also used by VFE as indicated in the following table by a mark of [Used by VFE]. As shown, inline data starts at dword 4 of MEDIA_OBJECT_EX command.

Inline data in AVC-IT mode ([DevCTG])

| DWord | Bit | Description |
|-------|-------|---|
| 4+0 | 31:27 | Reserved. MBZ |
| | 26:25 | <p>MbAffFieldFlag This field indicates that the current macroblock is a field macroblock within a MbAff frame picture. It is provided as Flag = MbaffFrame & FieldMbFlag. 00 = if (Flag == 0) 11 = if (Flag == 1) Other encodings are reserved</p> |



| DWord | Bit | Description |
|-------|-------|--|
| | 24 | <p>FieldMbPolarityFlag</p> <p>This field indicates the field polarity of the current macroblock.</p> <p>Within a MbAff frame picture, this field may be different per macroblock and is set to 1 only for the second macroblock in a MbAff pair if FieldMbFlag is set. Otherwise, it is set to 0.</p> <p>Within a field picture, this field is set to 1 if the current picture is the bottom field picture. Otherwise, it is set to 0. It is a constant for the whole field picture.</p> <p>This field is reserved and MBZ for a progressive frame picture.</p> <p>0 = Current macroblock is a field macroblock from the top field</p> <p>1 = Current macroblock is a field macroblock from the bottom field</p> <p><i>Programming Note: Here bits [26:24] (MbAffFieldFlag and FiedMbPolarityFlag) match with bits [10:8] of the Media Block Read message descriptor, simplifying the programming for message generation, as when MbAffFieldFlag is "1", kernels need to override the original "frame" surface state set for MBAFF frame picture.</i></p> |
| | 23 | Reserved: MBZ |
| | 22:20 | <p>MvSize (Motion Vector Size) [Used by VFE]. This field specifies the size of motion vectors for the macroblock stored in the indirect data buffer. The valid numbers are listed below indicating the size of the regrouped motion vectors. Details are provided in Section 1.7.3.1.4.</p> <p>This field is reserved (MBZ) when IntraMbFlag = 1.</p> <p>000 = 0: No motion vector</p> <p>001 reserved</p> <p>010 = 2MV: One motion vector pair</p> <p>011 reserved</p> <p>100 = 8MV: Four motion vector pairs</p> <p>101 = 16MV: 16 motion vectors</p> <p>110 = 32MV: 16 motion vector pairs</p> <p>111 reserved</p> |
| | 19 | DcBlockCodedYFlag [Used by VFE]. This field is consumed by VFE and is not delivered to the thread. |
| | 18 | DcBlockCodedCbFlag [Used by VFE]. This field is consumed by VFE and is not delivered to the thread. |
| | 17 | DcBlockCodedCrFlag [Used by VFE]. This field is consumed by VFE and is not delivered to the thread. |
| | 16 | Reserved. MBZ |
| | 15 | <p>Transform8x8Flag [Used by VFE]. This field equals to the value of <i>transform_size_8x8_flag</i> as defined in AVC spec. If set, it indicates that luma samples are in residual 8x8 blocks. Otherwise, it indicates that luma samples are in residual 4x4 blocks.</p> <p>0: luma residual 4x4 blocks</p> <p>1: luma residual 8x8 blocks</p> |
| | 14 | <p>FieldMbFlag (Field Macroblock Flag). This field specifies whether current macroblock is field macroblock.</p> <p>0 = Frame macroblock.</p> <p>1 = Field macroblock.</p> |



| DWord | Bit | Description |
|-------|------|--|
| | 13 | <p>IntraMbFlag (Intra Macroblock Flag). This field specifies whether the current macroblock is an Intra (I) macroblock. A collective macroblock type in AVC standard includes I, SI, P and B. Type SI is not supported.</p> <p>0 = P or B macroblock. 1 = I macroblock.</p> |
| | 12:8 | <p>MbType (Macroblock Type). This field, along with IntraMbFlag specifies the macroblock types.</p> <p>Further details can be found in Section 1.7.3.1.2.</p> |
| | 7:6 | <p>WeightedBiPredFlag (Weighted Bidirectional Prediction Flag) (from Picture State).</p> <p>This field specifies the bidirectional prediction mode and is derived from syntax elements <i>weighted_bipred_flag</i> and <i>weighted_pred_flag</i> as defined in AVC spec.</p> <p>It is valid only for inter predicted macroblock. Otherwise (intra macroblock), this field is reserved and MBZ.</p> <p>For B-macroblock, this field is the same as <i>weighted_bipred_flag</i> as defined in AVC spec.</p> <p>00 = Default weighted prediction 01 = Explicit weighted prediction 10 = Implicit weighted prediction 11 = Reserved.</p> <p>For P-macroblock, the MSB is always 0 and the LSB is the same as <i>weighted_pred_flag</i> as defined in AVC spec.</p> <p>00 = Default weighted prediction 01 = Explicit weighted prediction 10 and 11 are reserved</p> |
| | 5 | <p>WeightedPredFlag (Weighted Prediction Flag) [from Picture State].</p> <p>It is valid only for inter predicted macroblock. Otherwise (intra macroblock), this field is reserved and MBZ.</p> <p>0 = Default weighted prediction 1 = Explicit weighted prediction</p> <p><i>Note: Information in this field is also carried in WeightedBiPredFlag.</i></p> |
| | 4 | <p>IntraPredAvailFlagF – F (Pixel [-1, 7] available for intra prediction)</p> <p><i>F = Is_Left_MB_Field & Is_Left_Bottom_MB_Intra</i></p> |
| | 3:2 | <p>ChromaFormatIdc (Chroma Format Indicator). This field is equal to the value of <i>chroma_format_idct</i> as defined in AVC (§7.4.2.1). It specifies the chroma sampling relative to the luma sampling.</p> <p>This field is constant within a picture.</p> <p>00 = Luma only (monochrome) 01 = YUV420 sampling 10 is reserved (for YUV422 sampling) 11 is reserved (for YUV444 sampling)</p> |



| DWord | Bit | Description |
|------------------|-------|--|
| | 1 | <p>MbaffFrameFlag (MB-AFF Frame Flag) [PPS]. This field indicates whether the current picture is a progressive frame or a MB-AFF (macroblock adaptive frame field) frame picture.</p> <p>This field is frame Picture Parameter Set.</p> <p>This field is reserved (MBZ) if FieldPicFlag = 1.</p> <p>0 = Progressive frame picture 1 = MB-AFF frame picture</p> |
| | 0 | <p>FieldPicFlag (Field Picture Flag) [PPS]. This field indicates whether the current picture is a field or a frame picture. A frame picture may be a progressive frame picture or an MB-AFF frame picture depending on the value of MbaffFrame.</p> <p>This field is frame Picture Parameter Set.</p> <p>0 = Frame picture 1 = Field picture</p> |
| 4+1 | 31:16 | CbpY (Coded Block Pattern Y) [Used by VFE] |
| | 15:8 | <p>VertOrigin (Vertical Origin). This field specifies the vertical origin of current macroblock in the destination picture in units of macroblocks. For field macroblock pair in MBAFF frame, the vertical origins for both macroblocks should be set as if they were located in corresponding field pictures. For example, for field macroblock pair originated at (16, 64) pixel location in an MBAFF frame picture, the Vertical Origin for both macroblocks should be set as 2 (macroblocks).</p> <p>Format = U8 in unit of macroblock.</p> |
| | 7:0 | <p>HorzOrigin (Horizontal Origin). This field specifies the horizontal origin of current macroblock in the destination picture in units of macroblocks.</p> <p>Format = U8 in unit of macroblock.</p> |
| 4+2 | 31:24 | QpPrimeCr [Used by VFE] |
| | 23:16 | QpPrimeCb [Used by VFE] |
| | 15:8 | QpPrimeY [Used by VFE] |
| | 7:4 | CbpCr [Used by VFE] |
| | 3:0 | CbpCb [Used by VFE] |
| 4+3 to 4+5 | 31:0 | <p>For intra macroblocks, see Error! Reference source not found.</p> <p>For intra macroblocks, see</p> |
| 4+6 | 31:16 | <p>LevelScaleCb. [Used by VFE]</p> <p>This field is for inverse transform of the Chroma (Cb) DC block.</p> <p>The LevelScale field is consumed by VFE and is not needed by the thread, but since the GRF has to be filled to the end of the block it should be sent anyways.</p> |
| | 15:0 | <p>LevelScaleCr. [Used by VFE]</p> <p>This field is for inverse transform of the Chroma (Cr) DC block.</p> |
| 4+7 | 31:16 | Reserved . MBZ |
| | 15:0 | <p>LevelScaleY. [Used by VFE]</p> <p>This field is for inverse transform of the Luma DC block.</p> |



Inline data in AVC-IT mode ([DevILK])

| DWord | Bit | Description |
|-------|----------------------|--|
| 4+0 | 31:27 | Reserved. MBZ |
| | 26:25 | <p>MbAffFieldFlag</p> <p>This field indicates that the current macroblock is a field macroblock within a MbAff frame picture. It is provided as Flag = MbaffFrame & FieldMbFlag.</p> <p>00 = if (Flag == 0) 11 = if (Flag == 1)</p> <p>Other encodings are reserved</p> |
| | 24 | <p>FieldMbPolarityFlag</p> <p>This field indicates the field polarity of the current macroblock.</p> <p>Within a MbAff frame picture, this field may be different per macroblock and is set to 1 only for the second macroblock in a MbAff pair if FieldMbFlag is set. Otherwise, it is set to 0.</p> <p>Within a field picture, this field is set to 1 if the current picture is the bottom field picture. Otherwise, it is set to 0. It is a constant for the whole field picture.</p> <p>This field is reserved and MBZ for a progressive frame picture.</p> <p>0 = Current macroblock is a field macroblock from the top field 1 = Current macroblock is a field macroblock from the bottom field</p> <p><i>Programming Note: Here bits [26:24] (MbAffFieldFlag and FieldMbPolarityFlag) match with bits [10:8] of the Media Block Read message descriptor, simplifying the programming for message generation, as when MbAffFieldFlag is "1", kernels need to override the original "frame" surface state set for MBAFF frame picture.</i></p> |
| | 23 | Reserved: MBZ |
| | 22:20 | <p>MvSize (Motion Vector Size) [Used by VFE]. This field specifies the size of motion vectors for the macroblock stored in the indirect data buffer. The valid numbers are listed below indicating the size of the regrouped motion vectors. Details are provided in Section 1.7.3.1.4.</p> <p>This field is reserved (MBZ) when IntraMbFlag = 1.</p> <p>000 = 0: No motion vector 001 reserved 010 = 2MV: One motion vector pair 011 reserved 100 = 8MV: Four motion vector pairs 101 = 16MV: 16 motion vectors 110 = 32MV: 16 motion vector pairs 111 reserved</p> |
| | 19 | DcBlockCodedYFlag [Used by VFE]. This field is consumed by VFE and is not delivered to the thread. |
| | 18 | DcBlockCodedCbFlag [Used by VFE]. This field is consumed by VFE and is not delivered to the thread. |
| | 17 | DcBlockCodedCrFlag [Used by VFE]. This field is consumed by VFE and is not delivered to the thread. |
| 16 | Reserved. MBZ | |



| DWord | Bit | Description |
|-------|------|--|
| | 15 | Transform8x8Flag [Used by VFE]. This field equals to the value of <i>transform_size_8x8_flag</i> as defined in AVC spec. If set, it indicates that luma samples are in residual 8x8 blocks. Otherwise, it indicates that luma samples are in residual 4x4 blocks. 0: luma residual 4x4 blocks 1: luma residual 8x8 blocks |
| | 14 | FieldMbFlag (Field Macroblock Flag) . This field specifies whether current macroblock is field macroblock. 0 = Frame macroblock. 1 = Field macroblock. |
| | 13 | IntraMbFlag (Intra Macroblock Flag) . This field specifies whether the current macroblock is an Intra (I) macroblock. A collective macroblock type in AVC standard includes I, SI, P and B. Type SI is not supported. 0 = P or B macroblock. 1 = I macroblock. |
| | 12:8 | MbType (Macroblock Type) . This field, along with IntraMbFlag specifies the macroblock types. Further details can be found in Section 1.7.3.1.2. |
| | 7:6 | WeightedBiPredFlag (Weighted Bidirectional Prediction Flag) (from Picture State). This field specifies the bidirectional prediction mode and is derived from syntax elements <i>weighted_bipred_flag</i> and <i>weighted_pred_flag</i> as defined in AVC spec. It is valid only for inter predicted macroblock. Otherwise (intra macroblock), this field is reserved and MBZ. For B-macroblock, this field is the same as <i>weighted_bipred_flag</i> as defined in AVC spec. 00 = Default weighted prediction 01 = Explicit weighted prediction 10 = Implicit weighted prediction 11 = Reserved. For P-macroblock, the MSB is always 0 and the LSB is the same as <i>weighted_pred_flag</i> as defined in AVC spec. 00 = Default weighted prediction 01 = Explicit weighted prediction 10 and 11 are reserved |
| | 5 | WeightedPredFlag (Weighted Prediction Flag) [from Picture State]. It is valid only for inter predicted macroblock. Otherwise (intra macroblock), this field is reserved and MBZ. 0 = Default weighted prediction 1 = Explicit weighted prediction <i>Note: Information in this field is also carried in WeightedBiPredFlag.</i> |
| | 4 | Reserved. MBZ |



| DWord | Bit | Description |
|------------------|-------|--|
| | 3:2 | <p>ChromaFormatIdc (Chroma Format Indicator). This field is equal to the value of <i>chroma_format_idct</i> as defined in AVC (§7.4.2.1). It specifies the chroma sampling relative to the luma sampling.</p> <p>This field is constant within a picture.</p> <p>00 = Luma only (monochrome)</p> <p>01 = YUV420 sampling</p> <p>10 is reserved (for YUV422 sampling)</p> <p>11 is reserved (for YUV444 sampling)</p> |
| | 1 | <p>MbaffFrameFlag (MB-AFF Frame Flag) [PPS]. This field indicates whether the current picture is a progressive frame or a MB-AFF (macroblock adaptive frame field) frame picture.</p> <p>This field is frame Picture Parameter Set.</p> <p>This field is reserved (MBZ) if FieldPicFlag = 1.</p> <p>0 = Progressive frame picture</p> <p>1 = MB-AFF frame picture</p> |
| | 0 | <p>FieldPicFlag (Field Picture Flag) [PPS]. This field indicates whether the current picture is a field or a frame picture. A frame picture may be a progressive frame picture or an MB-AFF frame picture depending on the value of MbaffFrame.</p> <p>This field is frame Picture Parameter Set.</p> <p>0 = Frame picture</p> <p>1 = Field picture</p> |
| 4+1 | 31:16 | CbpY (Coded Block Pattern Y) [Used by VFE] |
| | 15:8 | <p>VertOrigin (Vertical Origin). This field specifies the vertical origin of current macroblock in the destination picture in units of macroblocks. For field macroblock pair in MBAFF frame, the vertical origins for both macroblocks should be set as if they were located in corresponding field pictures. For example, for field macroblock pair originated at (16, 64) pixel location in an MBAFF frame picture, the Vertical Origin for both macroblocks should be set as 2 (macroblocks).</p> <p>Format = U8 in unit of macroblock.</p> |
| | 7:0 | <p>HorzOrigin (Horizontal Origin). This field specifies the horizontal origin of current macroblock in the destination picture in units of macroblocks.</p> <p>Format = U8 in unit of macroblock.</p> |
| 4+2 | 31:24 | QpPrimeCr [Used by VFE] |
| | 23:16 | QpPrimeCb [Used by VFE] |
| | 15:8 | QpPrimeY [Used by VFE] |
| | 7:4 | CbpCr [Used by VFE] |
| | 3:0 | CbpCb [Used by VFE] |
| 4+3 to 4+5 | 31:0 | <p>For intra macroblocks, see Error! Reference source not found.</p> <p>For intra macroblocks, see</p> |



| DWord | Bit | Description |
|-------|-------|--|
| 4+6 | 31:16 | LevelScaleCb. [Used by VFE] This field is for inverse transform of the Chroma (Cb) DC block. The LevelScale field is consumed by VFE and is not needed by the thread, but since the GRF has to be filled to the end of the block it should be sent anyways. |
| | 15:0 | LevelScaleCr. [Used by VFE] This field is for inverse transform of the Chroma (Cr) DC block. |
| 4+7 | 31:16 | Reserved. MBZ |
| | 15:0 | LevelScaleY. [Used by VFE] This field is for inverse transform of the Luma DC block. |
| 4+8 | 31:0 | BSD_Pass_Thru_Data(Reserved to 0 in BSD mode) |
| 4+9 | 31:0 | BSD_Pass_Thru_Data(Reserved to 0 in BSD mode) |
| 4+10 | 31:0 | BSD_Pass_Thru_Data(Reserved to 0 in BSD mode) |
| 4+11 | 31:0 | BSD_Pass_Thru_Data(Reserved to 0 in BSD mode) |

This table shows dwords 4+3, 4+4 and 4+5 of the inline data. Luma intra prediction mode (LumaIntraPredModes) is provided as a fixed-size data structure. Details can be found in Section 0.

Inline data subfields for an Intra Macroblock in AVC-IT mode (and AVC-MC mode)

| Dword | Bit | Description | | | | | | | | | | | | | | | |
|-------|---|--|------|---------------------------------|-----|----------------------------|---|---|---|--|---|---|---|--|---|---|---|
| 4+3 | 31:8 | Reserved | | | | | | | | | | | | | | | |
| | 7:0 | MbIntraStruct (Macroblock Intra Structure) <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>MotionVerticalFieldSelect Index</th> </tr> </thead> <tbody> <tr> <td>7:6</td> <td>ChromaIntraPredMode</td> </tr> <tr> <td>5</td> <td>[DevILK] IntraPredAvailFlagF – F (Pixel [-1, 7] available for intra prediction) F = Is_Left_MB_Field & Is_Left_Bottom_MB_Intra [DevCTG] Reserved</td> </tr> <tr> <td>4</td> <td>IntraPredAvailFlagE – E (left neighbor bottom half)</td> </tr> <tr> <td>3</td> <td>IntraPredAvailFlagD – D (Upper right neighbor)</td> </tr> <tr> <td>2</td> <td>IntraPredAvailFlagC – C (Upper left neighbor)</td> </tr> <tr> <td>1</td> <td>IntraPredAvailFlagB – B (Upper neighbor)</td> </tr> <tr> <td>0</td> <td>IntraPredAvailFlagA – A (Left neighbor top half)</td> </tr> </tbody> </table> | Bits | MotionVerticalFieldSelect Index | 7:6 | ChromaIntraPredMode | 5 | [DevILK] IntraPredAvailFlagF – F (Pixel [-1, 7] available for intra prediction) F = Is_Left_MB_Field & Is_Left_Bottom_MB_Intra [DevCTG] Reserved | 4 | IntraPredAvailFlagE – E (left neighbor bottom half) | 3 | IntraPredAvailFlagD – D (Upper right neighbor) | 2 | IntraPredAvailFlagC – C (Upper left neighbor) | 1 | IntraPredAvailFlagB – B (Upper neighbor) | 0 |
| Bits | MotionVerticalFieldSelect Index | | | | | | | | | | | | | | | | |
| 7:6 | ChromaIntraPredMode | | | | | | | | | | | | | | | | |
| 5 | [DevILK] IntraPredAvailFlagF – F (Pixel [-1, 7] available for intra prediction) F = Is_Left_MB_Field & Is_Left_Bottom_MB_Intra [DevCTG] Reserved | | | | | | | | | | | | | | | | |
| 4 | IntraPredAvailFlagE – E (left neighbor bottom half) | | | | | | | | | | | | | | | | |
| 3 | IntraPredAvailFlagD – D (Upper right neighbor) | | | | | | | | | | | | | | | | |
| 2 | IntraPredAvailFlagC – C (Upper left neighbor) | | | | | | | | | | | | | | | | |
| 1 | IntraPredAvailFlagB – B (Upper neighbor) | | | | | | | | | | | | | | | | |
| 0 | IntraPredAvailFlagA – A (Left neighbor top half) | | | | | | | | | | | | | | | | |
| 4+4 | 31:16 | LumaIntraPredModes[1] (Luma Intra Prediction Modes) | | | | | | | | | | | | | | | |
| | 15:0 | LumaIntraPredModes[0] | | | | | | | | | | | | | | | |
| 4+5 | 31:16 | LumaIntraPredModes[3] | | | | | | | | | | | | | | | |
| | 15:0 | LumaIntraPredModes[2] | | | | | | | | | | | | | | | |

Dwords 4+3, 4+4 and 4+5 of the inline data for an inter-predicted macroblock is detailed in



Inline data subfields for an Inter Macroblock in AVC-IT mode (and AVC-MC mode)

| DWord | Bit | Description |
|--------------|------------|--|
| 4+3 | 31:24 | Log2WeightDenomChroma |
| | 23:16 | Log2WeightDenomLuma |
| | 15:8 | <p>SubMbPredMode (Sub Macroblock Prediction Mode)</p> <p>This field describes the prediction mode of the sub macroblocks. It contains four subfields each with 2-bits, corresponding to the 4 fixed size 8x8 sub macroblocks in sequential order. Details can be found in Section 1.7.3.1.3.</p> <p>This field is derived from sub_mb_type for a BP_8x8 macroblock.</p> <p>This field is derived from MbType for a non-BP_8x8 inter macroblock, and carries redundant information as MbType</p> <p>Bits [1:0]: SubMbPredMode[0] Bits [3:2]: SubMbPredMode[1] Bits [5:4]: SubMbPredMode[2] Bits [7:6]: SubMbPredMode[3]</p> |
| | 7:0 | <p>SubMbShape (Sub Macroblock Shape)</p> <p>This field describes the subdivision of the sub macroblocks. It contains four subfields each with 2-bits, corresponding to the 4 fixed size 8x8 sub macroblocks in sequential order. Details can be found in Section 1.7.3.1.3.</p> <p>This field is derived from sub_mb_type for a BP_8x8 macroblock.</p> <p>This field is forced to 0 for a non-BP_8x8 inter macroblock, and effectively carries redundant information as MbType.</p> <p>Bits [1:0]: SubMbShape[0] Bits [3:2]: SubMbShape[1] Bits [5:4]: SubMbShape[2] Bits [7:6]: SubMbShape[3]</p> |
| 4+4 | 31:24 | BindingTableIndexForward – Block 3 |
| | 23:16 | BindingTableIndexForward – Block 2 |
| | 15:8 | BindingTableIndexForward – Block 1 |
| | 7:0 | BindingTableIndexForward – Block 0 |
| 4+5 | 31:24 | BindingTableIndexBackward – Block 3 |
| | 23:16 | BindingTableIndexBackward – Block 2 |
| | 15:8 | BindingTableIndexBackward – Block 1 |
| | 7:0 | BindingTableIndexBackward – Block 0 |



1.7.3.1.1 Luma Intra Prediction Modes

Luma Intra Prediction Modes (LumaIntraPredModes) is defined in this table. It is further categorized as Intra16x16PredMode (**Error! Reference source not found.**), Intra8x8PredMode (**Error! Reference source not found.**) and Intra4x4PredMode (**Error! Reference source not found.**), operating on 16x16, 8x8 and 4x4 block sizes, respectively. **Error! Reference source not found.** illustrates the intra prediction directions geometrically for the Intra4x4 prediction. When a macroblock is subdivided, the intra prediction is performed for the subdivision in a predetermined order. For example, the figure below shows the block order for Intra4x4 prediction. And the figure below shows the block order of Block8x8 in a 16x16 region or Block4x4 in a 8x8 region.

Definition of LumaIntraPredModes

| LumaIntraPredModes [index] | | Intra16x16PredMode | Intra8x8PredMode | Intra4x4PredMode |
|-------------------------------|-------|---|------------------------------------|------------------------------------|
| Index | Bit | MbType = [1...24] Transform8x8Flag = 0 | MbType = 0 Transform8x8Flag = 1 | MbType = 0 Transform8x8Flag = 0 |
| 0 | 15:12 | MBZ | Block8x8 3 | Block4x4 3 (0_0) |
| | 11:8 | MBZ | Block8x8 2 | Block4x4 2 (0_1) |
| | 7:4 | MBZ | Block8x8 1 | Block4x4 1 (0_2) |
| | 3:0 | Block16x16 | Block8x8 0 | Block4x4 0 (0_3) |
| 1 | 15:12 | MBZ | MBZ | Block4x4 7 (1_0) |
| | 11:8 | MBZ | MBZ | Block4x4 6 (1_1) |
| | 7:4 | MBZ | MBZ | Block4x4 5 (1_2) |
| | 3:0 | MBZ | MBZ | Block4x4 4 (1_3) |
| 2 | 15:12 | MBZ | MBZ | Block4x4 11 (2_0) |
| | 11:8 | MBZ | MBZ | Block4x4 10 (2_1) |
| | 7:4 | MBZ | MBZ | Block4x4 9 (2_2) |
| | 3:0 | MBZ | MBZ | Block4x4 8 (2_3) |
| 3 | 15:12 | MBZ | MBZ | Block4x4 15 (3_0) |
| | 11:8 | MBZ | MBZ | Block4x4 14 (3_1) |
| | 7:4 | MBZ | MBZ | Block4x4 13 (3_2) |
| | 3:0 | MBZ | MBZ | Block4x4 12 (3_3) |



Definition of Intra16x16PredMode

| Intra16x16PredMode | Description |
|--------------------|------------------------|
| 0 | Intra_16x16_Vertical |
| 1 | Intra_16x16_Horizontal |
| 2 | Intra_16x16_DC |
| 3 | Intra_16x16_Plane |
| 4 – 15 | Reserved |


Definition of Intra8x8PredMode

| Intra8x8PredMode | Description |
|------------------|-------------------------------|
| 0 | Intra_8x8_Vertical |
| 1 | Intra_8x8_Horizontal |
| 2 | Intra_8x8_DC |
| 3 | Intra_8x8_Diagonal_Down_Left |
| 4 | Intra_8x8_Diagonal_Down_Right |
| 5 | Intra_8x8_Vertical_Right |
| 6 | Intra_8x8_Horizontal_Down |
| 7 | Intra_8x8_Vertical_Left |
| 8 | Intra_8x8_Horizontal_Up |
| 9 – 15 | Reserved |

Definition of Intra4x4PredMode

| Intra4x4PredMode | Description |
|------------------|-------------------------------|
| 0 | Intra_4x4_Vertical |
| 1 | Intra_4x4_Horizontal |
| 2 | Intra_4x4_DC |
| 3 | Intra_4x4_Diagonal_Down_Left |
| 4 | Intra_4x4_Diagonal_Down_Right |
| 5 | Intra_4x4_Vertical_Right |
| 6 | Intra_4x4_Horizontal_Down |
| 7 | Intra_4x4_Vertical_Left |
| 8 | Intra_4x4_Horizontal_Up |
| 9 – 15 | Reserved |

Intra_4x4 prediction mode directions



Numbers of Block4x4 in a 16x16 region

| | | | |
|----|----|----|----|
| 0 | 1 | 4 | 5 |
| 2 | 3 | 6 | 7 |
| 8 | 9 | 12 | 13 |
| 10 | 11 | 14 | 15 |



Numbers of Block4x4 in an 8x8 region or numbers of Block8x8 in a 16x16 region

| | |
|---|---|
| 0 | 1 |
| 2 | 3 |



1.7.3.1.2 Macroblock Type

Macroblock Type, MbType, is defined as a unified parameter for all slice types (I, P or B slices) as shown in **Error! Reference source not found.** Furthermore, MbType has the same meaning for a P macroblock and a B macroblock. For example, BP_L0_16x16 can be viewed as a P_L0_16x16 macroblock in a P slice or a B_L0_16x16 macroblock in a B slice.

As shown in **Error! Reference source not found.**, Macroblock Type (MbType) is derived from *mb_type*, as defined in AVC spec, for an I-, P- or B-slice.

Definition of MbType

| MbType | For an Intra Macroblock (IntraMbFlag = 1) | For an Inter Macroblock (IntraMbFlag = 1) |
|--------|--|--|
| 0 | I_NxN | Reserved |
| 1 | I_16x16_0_0_0 | BP_L0_16x16 |
| 2 | I_16x16_1_0_0 | B_L1_16x16 |
| 3 | I_16x16_2_0_0 | B_Bi_16x16 |
| 4 | I_16x16_3_0_0 | BP_L0_L0_16x8 |
| 5 | I_16x16_0_1_0 | BP_L0_L0_8x16 |
| 6 | I_16x16_1_1_0 | B_L1_L1_16x8 |
| 7 | I_16x16_2_1_0 | B_L1_L1_8x16 |
| 8 | I_16x16_3_1_0 | B_L0_L1_16x8 |
| 9 | I_16x16_0_2_0 | B_L0_L1_8x16 |
| 10 | I_16x16_1_2_0 | B_L1_L0_16x8 |
| 11 | I_16x16_2_2_0 | B_L1_L0_8x16 |
| 12 | I_16x16_3_2_0 | B_L0_Bi_16x8 |
| 13 | I_16x16_0_0_1 | B_L0_Bi_8x16 |
| 14 | I_16x16_1_0_1 | B_L1_Bi_16x8 |
| 15 | I_16x16_2_0_1 | B_L1_Bi_8x16 |
| 16 | I_16x16_3_0_1 | B_Bi_L0_16x8 |
| 17 | I_16x16_0_1_1 | B_Bi_L0_8x16 |
| 18 | I_16x16_1_1_1 | B_Bi_L1_16x8 |
| 19 | I_16x16_2_1_1 | B_Bi_L1_8x16 |
| 20 | I_16x16_3_1_1 | B_Bi_Bi_16x8 |
| 21 | I_16x16_0_2_1 | B_Bi_Bi_8x16 |
| 22 | I_16x16_1_2_1 | BP_8x8 |
| 23 | I_16x16_2_2_1 | Reserved |
| 24 | I_16x16_3_2_1 | Reserved |



| MbType | For an Intra Macroblock (IntraMbFlag = 1) | For an Inter Macroblock (IntraMbFlag = 1) |
|--------|--|--|
| 25 | I_PCM | Reserved |
| 26 | Reserved (for SI) | Reserved |
| 27-63 | Reserved | Reserved |

Deriving MbType from mb_type for I, P and B slices

| mb_type | I Slice | | P Slice | | B Slice | |
|---------|---------|---------------|---------|---------------|---------|---------------------------------|
| | MbType | Description | MbType | Description | MbType | Description |
| 0 | 0 | I_NxN | 1 | BP_L0_16x16 | 22 | B_Direct_16x16 mapped to BP_8x8 |
| 1 | 1 | I_16x16_0_0_0 | 4 | BP_L0_L0_16x8 | 1 | BP_L0_16x16 |
| 2 | 2 | I_16x16_1_0_0 | 5 | BP_L0_L0_8x16 | 2 | B_L1_16x16 |
| 3 | 3 | I_16x16_2_0_0 | 22 | BP_8x8 | 3 | B_Bi_16x16 |
| 4 | 4 | I_16x16_3_0_0 | 22 | BP_8x8 | 4 | BP_L0_L0_16x8 |
| 5 | 5 | I_16x16_0_1_0 | 0 | I_NxN | 5 | BP_L0_L0_8x16 |
| 6 | 6 | I_16x16_1_1_0 | 1 | I_16x16_0_0_0 | 6 | B_L1_L1_16x8 |
| 7 | 7 | I_16x16_2_1_0 | 2 | I_16x16_1_0_0 | 7 | B_L1_L1_8x16 |
| 8 | 8 | I_16x16_3_1_0 | 3 | I_16x16_2_0_0 | 8 | B_L0_L1_16x8 |
| 9 | 9 | I_16x16_0_2_0 | 4 | I_16x16_3_0_0 | 9 | B_L0_L1_8x16 |
| 10 | 10 | I_16x16_1_2_0 | 5 | I_16x16_0_1_0 | 10 | B_L1_L0_16x8 |
| 11 | 11 | I_16x16_2_2_0 | 6 | I_16x16_1_1_0 | 11 | B_L1_L0_8x16 |
| 12 | 12 | I_16x16_3_2_0 | 7 | I_16x16_2_1_0 | 12 | B_L0_Bi_16x8 |
| 13 | 13 | I_16x16_0_0_1 | 8 | I_16x16_3_1_0 | 13 | B_L0_Bi_8x16 |
| 14 | 14 | I_16x16_1_0_1 | 9 | I_16x16_0_2_0 | 14 | B_L1_Bi_16x8 |
| 15 | 15 | I_16x16_2_0_1 | 10 | I_16x16_2_0_1 | 15 | B_L1_Bi_8x16 |
| 16 | 16 | I_16x16_3_0_1 | 11 | I_16x16_2_2_0 | 16 | B_Bi_L0_16x8 |
| 17 | 17 | I_16x16_0_1_1 | 12 | I_16x16_3_2_0 | 17 | B_Bi_L0_8x16 |
| 18 | 18 | I_16x16_1_1_1 | 13 | I_16x16_0_0_1 | 18 | B_Bi_L1_16x8 |
| 19 | 19 | I_16x16_2_1_1 | 14 | I_16x16_1_0_1 | 19 | B_Bi_L1_8x16 |
| 20 | 20 | I_16x16_3_1_1 | 15 | I_16x16_2_0_1 | 20 | B_Bi_Bi_16x8 |
| 21 | 21 | I_16x16_0_2_1 | 16 | I_16x16_3_0_1 | 21 | B_Bi_Bi_8x16 |
| 22 | 22 | I_16x16_1_2_1 | 17 | I_16x16_0_1_1 | 22 | BP_8x8 |
| 23 | 23 | I_16x16_2_2_1 | 18 | I_16x16_1_1_1 | 0 | I_NxN |



| mb_type | I Slice | | P Slice | | B Slice | |
|---------|---------|---------------|---------|---------------|---------|---------------|
| | MbType | Description | MbType | Description | MbType | Description |
| 24 | 24 | I_16x16_3_2_1 | 19 | I_16x16_2_1_1 | 1 | I_16x16_0_0_0 |
| 25 | 25 | I_PCM | 20 | I_16x16_3_1_1 | 2 | I_16x16_1_0_0 |
| 26 | n/a | n/a | 21 | I_16x16_0_2_1 | 3 | I_16x16_2_0_0 |
| 27 | | | 22 | I_16x16_1_2_1 | 4 | I_16x16_3_0_0 |
| 28 | | | 23 | I_16x16_2_2_1 | 5 | I_16x16_0_1_0 |
| 29 | | | 24 | I_16x16_3_2_1 | 6 | I_16x16_1_1_0 |
| 30 | | | 25 | I_PCM | 7 | I_16x16_2_1_0 |
| 31 | | | n/a | n/a | 8 | I_16x16_3_1_0 |
| 32 | | | | | 9 | I_16x16_0_2_0 |
| 33 | | | | | 10 | I_16x16_2_0_1 |
| 34 | | | | | 11 | I_16x16_2_2_0 |
| 35 | | | | | 12 | I_16x16_3_2_0 |
| 36 | | | | | 13 | I_16x16_0_0_1 |
| 37 | | | | | 14 | I_16x16_1_0_1 |
| 38 | | | | | 15 | I_16x16_2_0_1 |
| 39 | | | | | 16 | I_16x16_3_0_1 |
| 40 | | | | | 17 | I_16x16_0_1_1 |
| 41 | | | | | 18 | I_16x16_1_1_1 |
| 42 | | | | | 19 | I_16x16_2_1_1 |
| 43 | | | | | 20 | I_16x16_3_1_1 |
| 44 | | | | | 21 | I_16x16_0_2_1 |
| 45 | | | | | 22 | I_16x16_1_2_1 |
| 46 | | | | | 23 | I_16x16_2_2_1 |
| 47 | | | | | 24 | I_16x16_3_2_1 |
| 48 | | | | | 25 | I_PCM |
| 49-63 | | | | | n/a | n/a |



1.7.3.1.3 Sub Macroblock Shape and Sub Macroblock Prediction Mode

Sub Macroblock Shape, SubMbShape, describes the shape of the sub divisions of an 8x8 sub macroblock of a BP_8x8 macroblock. Sub Macroblock Prediction Mode, SubMbPredMode, indicates the prediction mode for the sub macroblock. They are defined in **Error! Reference source not found.** and **Error! Reference source not found.**. Both of these parameters can be derived from sub_mb_type field as defined in AVC spec according to **Error! Reference source not found.** and **Error! Reference source not found.**.

For a non-BP_8x8 inter macroblock (IntraMbFlag = 0), the sub macroblocks will be greater than and equal to 8x8. Both SubMbShape and SubMbPredMode must be filled to match with the MbType. In particular, SubMbShape is 0 and SubMbPredMode is determined based on MbType according to **Error! Reference source not found.**.

Definition of SubMbShape for an 8x8 region of a BP_8x8 macroblock

| SubMbShape | NumSubMbPart | SubMbPartWidth | SubMbPartHeight |
|------------|--------------|----------------|-----------------|
| 0 | 1 | 8 | 8 |
| 1 | 2 | 8 | 4 |
| 2 | 2 | 4 | 8 |
| 3 | 4 | 4 | 4 |

Definition of SubMbPredMode for an 8x8 region of a BP_8x8 macroblock

| SubMbPredMode | Description | Comments |
|---------------|-------------|-----------------|
| 0 | Pred_L0 | P_8x8 and B_8x8 |
| 1 | Pred_L1 | B_8x8 only |
| 2 | BiPred | B_8x8 only |
| 3 | Reserved | |

Mapping sub_mb_type to SubMbType and SubMbPredMode in P macroblocks (BP_8x8)

| sub_mb_type [i] | name | SubMb Prediction | SubMbPartWidth | SubMbPartHeight | SubMbShape [i] | SubMbPredMode [i] |
|-----------------|----------|------------------|----------------|-----------------|----------------|-------------------|
| 0 | P_L0_8x8 | Pred_L0 | 8 | 8 | 0 | 0 |
| 1 | P_L0_8x4 | Pred_L0 | 8 | 4 | 1 | 0 |
| 2 | P_L0_4x8 | Pred_L0 | 4 | 8 | 2 | 0 |
| 3 | P_L0_4x4 | Pred_L0 | 4 | 4 | 3 | 0 |
| Inferred | n/a | n/a | n/a | n/a | n/a | n/a |



Mapping sub_mb_type to SubMbType and SubMbPredModd in B macroblocks (BP_8x8)

| sub_mb_type [i] | name | SubMb Prediction | SubMbPartWidth | SubMbPartHeight | SubMbShape [i] | SubMbPredMode [i] |
|-----------------|--------------|------------------|----------------|-----------------|----------------|-------------------|
| 0 | B_Direct_8x8 | Direct | 4 | 4 | 3 | ? |
| 1 | B_L0_8x8 | Pred_L0 | 8 | 8 | 0 | 0 |
| 2 | B_L1_8x8 | Pred_L1 | 8 | 8 | 0 | 1 |
| 3 | B_Bi_8x8 | BiPred | 8 | 8 | 0 | 2 |
| 4 | B_L0_8x4 | Pred_L0 | 8 | 4 | 1 | 0 |
| 5 | B_L0_4x8 | Pred_L0 | 4 | 8 | 2 | 0 |
| 6 | B_L1_8x4 | Pred_L1 | 8 | 4 | 1 | 1 |
| 7 | B_L1_4x8 | Pred_L1 | 4 | 8 | 2 | 1 |
| 8 | B_Bi_8x4 | BiPred | 8 | 4 | 1 | 2 |
| 9 | B_Bi_4x8 | BiPred | 4 | 8 | 2 | 2 |
| 10 | B_L0_4x4 | Pred_L0 | 4 | 4 | 3 | 0 |
| 11 | B_L1_4x4 | Pred_L1 | 4 | 4 | 3 | 1 |
| 12 | B_Bi_4x4 | BiPred | 4 | 4 | 3 | 2 |
| inferred | mb_type | Direct | 4 | 4 | 3 | ? |

SubMbPredMode[] for non BP_8x8 macroblocks (when IntraMbFlag = 0)

| MbType | Name | SubMbPredMode[i] | | | |
|--------|----------------------|------------------|----------|----------|----------|
| | | i = 0 | i = 1 | i = 2 | i = 3 |
| 0 | Reserved | Reserved | Reserved | Reserved | Reserved |
| 1 | BP_L0_16x16 | 0 | 0 | 0 | 0 |
| 2 | B_L1_16x16 | 1 | 1 | 1 | 1 |
| 3 | B_Bi_16x16 | 2 | 2 | 2 | 2 |
| 4 | BP_L0_L0_16x8 | 0 | 0 | 0 | 0 |
| 5 | BP_L0_L0_8x16 | 0 | 0 | 0 | 0 |
| 6 | B_L1_L1_16x8 | 1 | 1 | 1 | 1 |
| 7 | B_L1_L1_8x16 | 1 | 1 | 1 | 1 |
| 8 | B_L0_L1_16x8 | 0 | 0 | 1 | 1 |
| 9 | B_L0_L1_8x16 | 0 | 1 | 0 | 1 |
| 10 | B_L1_L0_16x8 | 1 | 1 | 0 | 0 |
| 11 | B_L1_L0_8x16 | 1 | 0 | 1 | 0 |



| MbType | Name | SubMbPredMode[i] | | | |
|--------|--------------|------------------|-------|-------|-------|
| | | i = 0 | i = 1 | i = 2 | i = 3 |
| 12 | B_L0_Bi_16x8 | 0 | 0 | 2 | 2 |
| 13 | B_L0_Bi_8x16 | 0 | 2 | 0 | 2 |
| 14 | B_L1_Bi_16x8 | 1 | 1 | 2 | 2 |
| 15 | B_L1_Bi_8x16 | 1 | 2 | 1 | 2 |
| 16 | B_Bi_L0_16x8 | 2 | 2 | 0 | 0 |
| 17 | B_Bi_L0_8x16 | 2 | 0 | 2 | 0 |
| 18 | B_Bi_L1_16x8 | 2 | 2 | 1 | 1 |
| 19 | B_Bi_L1_8x16 | 2 | 1 | 2 | 1 |
| 20 | B_Bi_Bi_16x8 | 2 | 2 | 2 | 2 |
| 21 | B_Bi_Bi_8x16 | 2 | 2 | 2 | 2 |

1.7.3.1.4 Motion Vector Size

In AVC, a macroblock may have 0 or 32 motion vectors and many other combinations in between. In order to simplify the AVC-IT interface, the motion vectors of a macroblock are regrouped. As shown in the table below, only 5 distinct combined motion vector states (cMvState) B0, B1, B2, P3 and B3, are derived, corresponding the MvSize of 0, 2, 8, 16, and 32, respectively.

The maximum value of MvSize depends on the profile and level of the input AVC data. According to AVC Spec Table A-4 in section A.3.3.2, for Main and High Profiles at Level greater than 3.0, MinLumaBiPreSize is set to 8x8 (i.e. sub_mb_type in B macroblocks shall not be equal to B_Bi_8x4, B_Bi_4x8, or B_Bi_4x4). Therefore, B3 state is not valid for the given profile and level.

Programming Notes: Programmers may (and should) take advantage of such profile and level restriction to conserve memory foot print for indirect data, memory bandwidth for delivering data as well as possibly the GRF register space storing motion vectors. For example, when the maximum possible MvSize is 16, only 16 dwords need to be allocated for motion vectors in both indirect data buffer and GRF space.

Motion vector regroup

| Mblk Type | MV State | Max # MVs | Reference Lists | Combined MV State (cMvState) | Comments |
|-----------|----------|-----------|-----------------|------------------------------|---|
| P | P0 | 0 | n/a | B0 | Merged with B0 |
| P | P1 | 1 | L0 | B1 | Merged with B1 |
| P | P2 | 4 | L0 | B2 | Merged with B2 |
| P | P3 | 16 | L0 | P3 | Sub-macroblock partition smaller than 8x8 |



| Mblk Type | MV State | Max # MVs | Reference Lists | Combined MV State (cMvState) | Comments |
|-----------|----------|-----------|-----------------|------------------------------|---|
| B | B0 | 0 | n/a | B0 | |
| B | B1 | 2 | L0, L1, or Bi | B1 | |
| B | B2 | 8 | L0, L1, or Bi | B2 | Sub-macroblock partition down to 8x8 |
| B | B3 | 32 | L0, L1, or Bi | B3 | For a High Profile AVC data, only encountered with level <= 3.1 |

cMvState can be derived based on the following macroblock parameters: MbType, SubMbShape, and SubMbPredMode. **Error! Reference source not found.** provides the detailed mapping.

Regrouped motion vector states for an Inter Macroblock

| MbType | Inter Macroblock Type | Max (sub_mb_type[]) | Max (SubMBPredMode[]) | Extract MV # | MV State | MvSize | Comments |
|--------|-----------------------|---------------------|-----------------------|--------------|----------|--------|----------|
| 1 | BP_L0_16x16 | n/a | n/a | 1 | B1 | 2MV | |
| 2 | B_L1_16x16 | n/a | n/a | 1 | B1 | 2MV | |
| 3 | B_Bi_16x16 | n/a | n/a | 2 | B1 | 2MV | |
| 4 | BP_L0_L0_16x8 | n/a | n/a | 2 | B2 | 8MV | |
| 5 | BP_L0_L0_8x16 | n/a | n/a | 2 | B2 | 8MV | |
| 6 | B_L1_L1_16x8 | n/a | n/a | 2 | B2 | 8MV | |
| 7 | B_L1_L1_8x16 | n/a | n/a | 2 | B2 | 8MV | |
| 8 | B_L0_L1_16x8 | n/a | n/a | 2 | B2 | 8MV | |
| 9 | B_L0_L1_8x16 | n/a | n/a | 2 | B2 | 8MV | |
| 10 | B_L1_L0_16x8 | n/a | n/a | 2 | B2 | 8MV | |
| 11 | B_L1_L0_8x16 | n/a | n/a | 2 | B2 | 8MV | |
| 12 | B_L0_Bi_16x8 | n/a | n/a | 3 | B2 | 8MV | |
| 13 | B_L0_Bi_8x16 | n/a | n/a | 3 | B2 | 8MV | |
| 14 | B_L1_Bi_16x8 | n/a | n/a | 3 | B2 | 8MV | |
| 15 | B_L1_Bi_8x16 | n/a | n/a | 3 | B2 | 8MV | |
| 16 | B_Bi_L0_16x8 | n/a | n/a | 3 | B2 | 8MV | |
| 17 | B_Bi_L0_8x16 | n/a | n/a | 3 | B2 | 8MV | |



| MbType | Inter Macroblock Type | Max (sub_mb_type[]) | Max (SubMBPredModel[]) | Extract MV # | MV State | MVSize | Comments |
|--------|-----------------------|---------------------|------------------------|--------------|----------|--------|------------------------------------|
| 18 | B_Bi_L1_16x8 | n/a | n/a | 3 | B2 | 8MV | |
| 19 | B_Bi_L1_8x16 | n/a | n/a | 3 | B2 | 8MV | |
| 20 | B_Bi_Bi_16x8 | n/a | n/a | 4 | B2 | 8MV | |
| 21 | B_Bi_Bi_8x16 | n/a | n/a | 4 | B2 | 8MV | |
| 22 | BP_8x8 | 0 | 1 | 4 | B2 | 8MV | Without sub-partition, no BiPred |
| 22 | BP_8x8 | 0 | 2 | 5,6,7,8 | B2 | 8MV | Without sub-partition, with BiPred |
| 22 | BP_8x8 | > 0 | 1 | 5-16 | P3 | 16MV | With sub-partition, no BiPred |
| 22 | BP_8x8 | > 0 | 2 | 6-32 | B3 | 32MV | With sub-partition, with BiPred |

1.7.3.1.5 Binding Table Index Data in AVC-IT Mode

There are always 8 binding table indices transferred in the inline data for an Inter Macroblock, a forward and backward index for each 8x8 block in the macroblock. This data is derived from the reference index sent with each motion vector; since between 0 and 32 motion vectors can be sent, a mapping scheme is specified here to indicate which reference index is to be used for which block in the inline data.

The general scheme is that whenever the motion vectors are for partitions smaller than 8x8 then pick the upper right, since all binding table indices are guaranteed to be the same for all sub-blocks in an 8x8. If the motion vectors are for partitions larger than 8x8, then replicate the single binding table index for all 8x8s in the partition. If there is only a forward or backward motion vector specified, then replicate the binding table indices for the missing direction.

| MbType | Inter Macroblock | Binding Table Replication Rule |
|--------|----------------------|--|
| 1 | BP_L0_16x16 | L0 binding table index replicated to all 4 forward and all 4 backward |
| 2 | B_L1_16x16 | L1 binding table index replicated to all 4 forward and all 4 backward |
| 3 | B_Bi_16x16 | L0 replicated to all 4 forward, L1 replicated to all 4 backward |
| 4 | BP_L0_L0_16x8 | First L0 (top) replicated to blocks 0 & 1, both forward and backward, 2 nd L0 replicated to blocks 2 & 3, both forward and backward. |
| 5 | BP_L0_L0_8x16 | First L0 (left) replicated to blocks 0 & 2, both forward and backward, 2 nd L0 replicated to blocks 1 & 3, both forward and backward. |
| 6 | B_L1_L1_16x8 | First L1 (top) replicated to blocks 0 & 1, both forward and backward, 2 nd L1 replicated to blocks 2 & 3, both forward and backward. |



| MbType | Inter Macroblock | Binding Table Replication Rule |
|--------|------------------|---|
| 7 | B_L1_L1_8x16 | First L1 (left) replicated to blocks 0 & 2, both forward and backward, 2 nd L1 replicated to blocks 1 & 3, both forward and backward. |
| 8 | B_L0_L1_16x8 | First L0 (top) replicated to blocks 0 & 1, both forward and backward, 2 nd L1 replicated to blocks 2 & 3, both forward and backward. |
| 9 | B_L0_L1_8x16 | First L0 (left) replicated to blocks 0 & 2, both forward and backward, 2 nd L1 replicated to blocks 1 & 3, both forward and backward. |
| 10 | B_L1_L0_16x8 | First L1 (top) replicated to blocks 0 & 1, both forward and backward, 2 nd L0 replicated to blocks 2 & 3, both forward and backward. |
| 11 | B_L1_L0_8x16 | First L1 (left) replicated to blocks 0 & 2, both forward and backward, 2 nd L0 replicated to blocks 1 & 3, both forward and backward. |
| 12 | B_L0_Bi_16x8 | First L0 (top) replicated to blocks 0 & 1, both forward and backward, 2 nd L0 replicated to blocks forward 2 & 3, 2 nd L1 to backward blocks 2 & 3 |
| 13 | B_L0_Bi_8x16 | First L0 (left) replicated to blocks 0 & 2, both forward and backward, 2 nd L0 replicated to blocks forward 1 & 3, 2 nd L1 to backward blocks 2 & 3 |
| 14 | B_L1_Bi_16x8 | First L1 (top) replicated to blocks 0 & 1, both forward and backward, 2 nd L0 replicated to blocks forward 2 & 3, 2 nd L1 to backward blocks 2 & 3 |
| 15 | B_L1_Bi_8x16 | First L1 (left) replicated to blocks 0 & 2, both forward and backward, 2 nd L0 replicated to blocks forward 1 & 3, 2 nd L1 to backward blocks 2 & 3 |
| 16 | B_Bi_L0_16x8 | First L0 replicated to blocks forward 0 & 1, 1 st L1 to backward blocks 0 & 1, 2 nd L0 replicated to blocks 2 & 3, both forward and backward |
| 17 | B_Bi_L0_8x16 | First L0 replicated to blocks forward 0 & 2, 1 st L1 to backward blocks 0 & 2, 2 nd L0 replicated to blocks 1 & 3, both forward and backward |
| 18 | B_Bi_L1_16x8 | First L0 replicated to blocks forward 0 & 1, 1 st L1 to backward blocks 0 & 1, 2 nd L1 replicated to blocks 2 & 3, both forward and backward |
| 19 | B_Bi_L1_8x16 | First L0 replicated to blocks forward 0 & 2, 1 st L1 to backward blocks 0 & 2, 2 nd L1 replicated to blocks 1 & 3, both forward and backward |
| 20 | B_Bi_Bi_16x8 | First L0 replicated to forward blocks 0 & 1, 1 st L1 to backward blocks 0 & 1, 2 nd L0 replicated to forward blocks 2 & 3, 2 nd L1 to backward blocks 2 & 3 |
| 21 | B_Bi_Bi_8x16 | First L0 replicated to forward blocks 0 & 2, 1 st L1 to backward blocks 0 & 2, 2 nd L0 replicated to forward blocks 2 & 3, 2 nd L1 to backward blocks 2 & 3 |
| 22 | BP_8x8 | <p>1) If the MvSize is 8, then the Binding Table Indices can be directly derived from the reference indices in the 8 motion vectors.</p> <p>2) If MvSize is 16, then the macroblock is being split into 4x4 sub-blocks and biprediction is off (only 1 motion vector per 4x4). In this case, each 4x4 can either be forward or backward predicted, but the table reference for each set of 4 in an 8x8 is the same. Each of the 4 motion vectors in an 8x8 needs to be looked at – if one of them is forward predicted then the associated table reference can be used for that 8x8 block, and if one is backward predicted then that can be used for the backward reference for the 8x8. If all 4 motion vectors are forward, then the backward reference is not used and the forward table reference can be used as the default.</p> <p>3) If MvSize is 32, then BiPred for the 4x4 sub-blocks. In this case between 4 and 8 motion vectors are sent per 8x8 block depending on whether the prediction is Bi or forward or backward. These motion vectors have to be searched in a similar way to the MvSize=16 case to find both the forward and backward reference or to replicate the existing reference if one of them is missing entirely.</p> |



1.7.3.2 Indirect Data Format in AVC-IT Mode

Indirect data in AVC-IT mode consist of Motion Vectors, Weight/Offset and Transform-domain Residue (Coefficient). All three data blocks have variable size. Sizes of Motion Vector block and the Weight-Offset block are determined by the MvSize value as shown in **Error! Reference source not found.** Weight-Offset block, if present, is always packed behind the Motion Vector block. Coefficient data block can be either packed behind the Weight-Offset block or start at a predetermined offset, controlled by the fields in VFE_STATE_EX.

When coefficient data block is packed behind, it starts at the next 8-dword aligned offset from the indirect object data address. This 8-dword alignment doesn't leave any gap between the coefficient data block from the motion vector data block and weight-offset data block with one exception. When MvSize = 2 and weight-offset is not present, there is a 4-dword gap. Hardware ignores the value in the gap.

Indirect subfield size in AVC-IT mode (and AVC-MC mode)

| MvSize | MV | | Weight/Offset | | Examples |
|--------|-------|----|---------------|----|--|
| | Count | DW | Count | DW | |
| 0 | 0 | 0 | 0 | 0 | Intra macroblock in a picture containing P and/or B slices |
| 2 | 2 | 4 | 1 | 4 | P or B macroblocks with 16x16 sub macroblock |
| 8 | 8 | 8 | 4 | 16 | P or B macroblocks with minimal sub macroblock at 8x8 |
| 16 | 16 | 16 | 4 | 16 | P macroblocks with minimal sub macroblock at less than 8x8 |
| 32 | 32 | 32 | 4 | 16 | B macroblocks with minimal sub macroblock at less than 8x8 |

1.7.3.2.1 Motion Vector Block of Indirect Data in AVC-IT and AVC-MC Modes

Motion Vector block contains motion vectors in an intermediate format that is partially expanded according to the smallest subdivisions within an inter-predicted macroblock. During the expansion (done by AVC BSD engine or done by host software), a place that does not contain a motion vector is filled by replicating the most relevant motion vector according to the following motion vector replication rules. The intent of such motion vector replication is to allow a simpler kernel programming with fewer conditions to check. This would likely reduce the kernel footprint; however, it may or may not achieve better performance.

Motion Vector Replication Rules:

- Rule #1
 - #1.1: For L0 MV, for any partition or subpartition where there is at least one motion vector
 - If L0 inter prediction exists, the corresponding L0 MV is used
 - Else if L1 inter prediction exists (of the same block), set to the same as L1 MV
 - (Note that there is no 'else' here. If the partition or subpartition doesnot contain a motion vector, it will be filled according to the following replication rules)
 - #1.2: For L1 MV, for any partition or subpartition where there is at least one motion vector
 - If L1 inter prediction exists, the corresponding L1 MV is used



- Else if L0 inter prediction exits (of the same block), set to the same as L0 MV
 - (Note that there is no ‘else’ here. If the partition or subpartition doesnot contain a motion vector, it will be filled according to the following replication rules)
- For a 16x16 partitioned macroblocked, MvSize = 2. The two MV fields follow Rule #1.
- For a macroblock with partition down to 8x8, MvSize = 8. The eight MV fields follow Rule #1.
 - For an 8x16 partition, each 8x16 is broken down into 2 8x8 stacking vertically. The 8x16 MVs (after rule #1) are replicated into both 8x8 blocks.
 - For an 16x8 partition, each 16x8 is broken down into 2 8x8 stacking horizontally. The 16x8 MVs (after rule #1) are replicated into both 8x8 blocks.
 - For an 8x8 partition, each 8x8 has its own MVs (after rule #1).
- For P macroblock with subpartition below 8x8, MvSize = 16,
 - For an 8x8 partition, the 8x8 L0 MV is replicated into all the four 4x4 blocks.
 - For an 4x8 subpartition within an 8x8 partition, each 4x8 is broken down into 2 4x4 stacking vertically. The 4x8 L0 MV is replicated into both 4x4 blocks.
 - For an 8x4 subpartition within an 8x8 partition, each 8x4 is broken down into 2 4x4 stacking horizontally. The 8x4 MV is replicated into both 4x4 blocks.
 - For a 4x4 subpartition within an 8x8 partition, each 4x4 has its own L0 MV.
- For B macroblock with subpartition below 8x8, MvSize = 32,
 - For an 8x8 partition, the 8x8 MVs (after rule #1) is replicated into all the four 4x4 blocks.
 - For an 4x8 subpartition within an 8x8 partition, each 4x8 is broken down into 2 4x4 stacking vertically. The 4x8 MVs (after rule #1) are replicated into both 4x4 blocks.
 - For an 8x4 subpartition within an 8x8 partition, each 8x4 is broken down into 2 4x4 stacking horizontally. The 8x4 MVs (after rule #1) are replicated into both 4x4 blocks.
 - For a 4x4 subpartition within an 8x8 partition, each 4x4 has its own MVs (after rule #1).

Indirect data Motion Vector block in AVC-IT mode (and AVC-MC mode)

| DWord | Bit | MvSize | | | | |
|-------|-------|--------|-----------|--------------|--------------|--------------|
| | | 0 | 2 | 8 | 16 | 32 |
| 0 | 31:16 | n/a | MVVert_L0 | MVVert_Y0_L0 | MVVert_Y0_L0 | MVVert_Y0_L0 |
| | 15:0 | n/a | MVHorz_L0 | MVHorz_Y0_L0 | MVHorz_Y0_L0 | MVHorz_Y0_L0 |
| 1 | 31:16 | n/a | MVVert_L1 | MVVert_Y0_L1 | MVVert_Y1_L0 | MVVert_Y0_L1 |
| | 15:0 | n/a | MVHorz_L1 | MVHorz_Y0_L1 | MVHorz_Y1_L0 | MVHorz_Y0_L1 |



| DWord | Bit | MvSize | | | | |
|-------|------|--------|---------------|----------|-----------|-----------|
| | | 0 | 2 | 8 | 16 | 32 |
| 2 | 31:0 | n/a | Reserved: MBZ | MV_Y1_L0 | MV_Y2_L0 | MV_Y1_L0 |
| 3 | 31:0 | n/a | Reserved: MBZ | MV_Y1_L1 | MV_Y3_L0 | MV_Y1_L1 |
| 4 | 31:0 | n/a | n/a | MV_Y2_L0 | MV_Y4_L0 | MV_Y2_L0 |
| 5 | 31:0 | n/a | n/a | MV_Y2_L1 | MV_Y5_L0 | MV_Y2_L1 |
| 6 | 31:0 | n/a | n/a | MV_Y3_L0 | MV_Y6_L0 | MV_Y3_L0 |
| 7 | 31:0 | n/a | n/a | MV_Y3_L1 | MV_Y7_L0 | MV_Y3_L1 |
| 8 | 31:0 | n/a | n/a | n/a | MV_Y8_L0 | MV_Y4_L0 |
| 9 | 31:0 | n/a | n/a | n/a | MV_Y9_L0 | MV_Y4_L1 |
| 10 | 31:0 | n/a | n/a | n/a | MV_Y10_L0 | MV_Y5_L0 |
| 11 | 31:0 | n/a | n/a | n/a | MV_Y11_L0 | MV_Y5_L1 |
| 12 | 31:0 | n/a | n/a | n/a | MV_Y12_L0 | MV_Y6_L0 |
| 13 | 31:0 | n/a | n/a | n/a | MV_Y13_L0 | MV_Y6_L1 |
| 14 | 31:0 | n/a | n/a | n/a | MV_Y14_L0 | MV_Y7_L0 |
| 15 | 31:0 | n/a | n/a | n/a | MV_Y15_L0 | MV_Y7_L1 |
| 16 | 31:0 | n/a | n/a | n/a | n/a | MV_Y8_L0 |
| 17 | 31:0 | n/a | n/a | n/a | n/a | MV_Y8_L1 |
| 18 | 31:0 | n/a | n/a | n/a | n/a | MV_Y9_L0 |
| 19 | 31:0 | n/a | n/a | n/a | n/a | MV_Y9_L1 |
| 20 | 31:0 | n/a | n/a | n/a | n/a | MV_Y10_L0 |
| 21 | 31:0 | n/a | n/a | n/a | n/a | MV_Y10_L1 |
| 22 | 31:0 | n/a | n/a | n/a | n/a | MV_Y11_L0 |
| 23 | 31:0 | n/a | n/a | n/a | n/a | MV_Y11_L1 |
| 24 | 31:0 | n/a | n/a | n/a | n/a | MV_Y12_L0 |
| 25 | 31:0 | n/a | n/a | n/a | n/a | MV_Y12_L1 |
| 26 | 31:0 | n/a | n/a | n/a | n/a | MV_Y13_L0 |
| 27 | 31:0 | n/a | n/a | n/a | n/a | MV_Y13_L1 |
| 28 | 31:0 | n/a | n/a | n/a | n/a | MV_Y14_L0 |
| 29 | 31:0 | n/a | n/a | n/a | n/a | MV_Y14_L1 |
| 30 | 31:0 | n/a | n/a | n/a | n/a | MV_Y15_L0 |
| 31 | 31:0 | n/a | n/a | n/a | n/a | MV_Y15_L1 |



1.7.3.2.2 Weight-Offset Block of Indirect Data in AVC-IT and AVC_MC Modes for WeightedBiPredFlag ≠ 10

Indirect data Weight-Offset block in AVC-IT mode (and AVC-MC mode)

Weight-Offset programming model[Errata – Dev Ctg/Dev EL] Driver needs to look ahead all slices for the frame and check if any weight value is 128. This can only happen when luma_weight_10/11/chroma_weight_10/11 flags are 0. If it finds a 128 then it needs to remap this value to a non used value within a 16 bit twos compliment range and then give the remapped number to the kernel for identifying this case.

Indirect data Weight-Offset block in AVC-IT mode (and AVC-MC mode) – [DevCTG]

| DWord | Bit | MvSize | | |
|-------|-------|--------|--|--|
| | | 0 | 2 | 8, 16, 32 |
| 0 | 31:24 | n/a | Offset_Y_L1 | Offset_Y_Block0_L1 |
| | 23:16 | n/a | Weight_Y_L1 | Weight_Y_Block0_L1 |
| | 15:8 | n/a | Offset_Y_L0 | Offset_Y_Block0_L0 |
| | 7:0 | n/a | Weight_Y_L0 | Weight_Y_Block0_L0 |
| 1 | 31:16 | n/a | WO_Cb_L1 | WO_Cb_Block0_L1 |
| | 15:0 | n/a | WO_Cb_L0 | WO_Cb_Block0_L0 |
| 2 | 31:16 | n/a | WO_Cr_L1 | WO_Cr_Block0_L1 |
| | 15:0 | n/a | WO_Cr_L0 | WO_Cr_Block0_L0 |
| 3 | 31:4 | n/a | Reserved: MBZ | Reserved: MBZ |
| | 3:0 | | [DevCTG] Reserved: MBZ [DevILK] Weight is 128 [ChromaL1, Chroma L0, Luma L1, Luma L0] | [DevCTG] Reserved: MBZ [DevILK] Weight is 128 [ChromaL1, Chroma L0, Luma L1, Luma L0] |
| 4 | 31:16 | n/a | n/a | WO_Y_Block1_L1 |
| | 15:0 | n/a | n/a | WO_Y_Block1_L0 |
| 5 | 31:16 | n/a | n/a | WO_Cb_Block1_L1 |
| | 15:0 | n/a | n/a | WO_Cb_Block1_L0 |
| 6 | 31:16 | n/a | n/a | WO_Cr_Block1_L1 |
| | 15:0 | n/a | n/a | WO_Cr_Block1_L0 |
| 7 | 31:4 | n/a | n/a | Reserved: MBZ |
| | 3:0 | | | [DevCTG] Reserved: MBZ [DevILK] Weight is 128[ChromaL1, Chroma L0, Luma L1, Luma L0] |
| 8 | 31:16 | n/a | n/a | WO_Y_Block2_L1 |
| | 15:0 | n/a | n/a | WO_Y_Block2_L0 |
| 9 | 31:16 | n/a | n/a | WO_Cb_Block2_L1 |
| | 15:0 | n/a | n/a | WO_Cb_Block2_L0 |



| DWord | Bit | MvSize | | |
|-------|-------|--------|-----|---|
| | | 0 | 2 | 8, 16, 32 |
| 10 | 31:16 | n/a | n/a | WO_Cr_Block2_L1 |
| | 15:0 | n/a | n/a | WO_Cr_Block2_L0 |
| 11 | 31:4 | n/a | n/a | Reserved: MBZ |
| | 3:0 | | | [DevCTG] Reserved: MBZ [DevILK] Weight is 128[ChromaL1, Chroma L0, Luma L1, Luma L0] |
| 12 | 31:16 | n/a | n/a | WO_Y_Block3_L1 |
| | 15:0 | n/a | n/a | WO_Y_Block3_L0 |
| 13 | 31:16 | n/a | n/a | WO_Cb_Block3_L1 |
| | 15:0 | n/a | n/a | WO_Cb_Block3_L0 |
| 14 | 31:16 | n/a | n/a | WO_Cr_Block3_L1 |
| | 15:0 | n/a | n/a | WO_Cr_Block3_L0 |
| 15 | 31:4 | n/a | n/a | Reserved: MBZ |
| | 3:0 | | | [DevCTG] Reserved: MBZ [DevILK] Weight is 128[ChromaL1, Chroma L0, Luma L1, Luma L0] |

[DevILK]

| DWord | Bit | MvSize | | |
|-------|-------|--------|--|--|
| | | 0 | 2 | 8, 16, 32 |
| 0 | 31:24 | n/a | Offset_Y_L1 | Offset_Y_Block0_L1 |
| | 23:16 | n/a | Weight_Y_L1 | Weight_Y_Block0_L1 |
| | 15:8 | n/a | Offset_Y_L0 | Offset_Y_Block0_L0 |
| | 7:0 | n/a | Weight_Y_L0 | Weight_Y_Block0_L0 |
| 1 | 31:16 | n/a | WO_Cb_L1 | WO_Cb_Block0_L1 |
| | 15:0 | n/a | WO_Cb_L0 | WO_Cb_Block0_L0 |
| 2 | 31:16 | n/a | WO_Cr_L1 | WO_Cr_Block0_L1 |
| | 15:0 | n/a | WO_Cr_L0 | WO_Cr_Block0_L0 |
| 3 | 31:4 | n/a | Reserved MBZ | Reserved MBZ |
| | 3:0 | n/a | Weight is 128[ChromaL1, Chroma L0, Luma L1, Luma L0] | Weight is 128[ChromaL1, Chroma L0, Luma L1, Luma L0] |
| 4 | 31:16 | n/a | n/a | WO_Y_Block1_L1 |
| | 15:0 | n/a | n/a | WO_Y_Block1_L0 |
| 5 | 31:16 | n/a | n/a | WO_Cb_Block1_L1 |



| DWord | Bit | MvSize | | |
|-------|-------|--------|-----|--|
| | | 0 | 2 | 8, 16, 32 |
| | 15:0 | n/a | n/a | WO_Cb_Block1_L0 |
| 6 | 31:16 | n/a | n/a | WO_Cr_Block1_L1 |
| | 15:0 | n/a | n/a | WO_Cr_Block1_L0 |
| 7 | 31:4 | n/a | n/a | Reserved MBZ |
| | 3:0 | n/a | n/a | Weight is 128[ChromaL1, Chroma L0, Luma L1, Luma L0] |
| 8 | 31:16 | n/a | n/a | WO_Y_Block2_L1 |
| | 15:0 | n/a | n/a | WO_Y_Block2_L0 |
| 9 | 31:16 | n/a | n/a | WO_Cb_Block2_L1 |
| | 15:0 | n/a | n/a | WO_Cb_Block2_L0 |
| 10 | 31:16 | n/a | n/a | WO_Cr_Block2_L1 |
| | 15:0 | n/a | n/a | WO_Cr_Block2_L0 |
| 11 | 31:4 | n/a | n/a | Reserved MBZ |
| | 3:0 | n/a | n/a | Weight is 128[ChromaL1, Chroma L0, Luma L1, Luma L0] |
| 12 | 31:16 | n/a | n/a | WO_Y_Block3_L1 |
| | 15:0 | n/a | n/a | WO_Y_Block3_L0 |
| 13 | 31:16 | n/a | n/a | WO_Cb_Block3_L1 |
| | 15:0 | n/a | n/a | WO_Cb_Block3_L0 |
| 14 | 31:16 | n/a | n/a | WO_Cr_Block3_L1 |
| | 15:0 | n/a | n/a | WO_Cr_Block3_L0 |
| 15 | 31:4 | n/a | n/a | Reserved MBZ |
| | 3:0 | n/a | n/a | Weight is 128[ChromaL1, Chroma L0, Luma L1, Luma L0] |



1.7.3.2.3 Weight-Offset Block of Indirect Data in AVC-IT and AVC_MC Modes for WeightedBiPredFlag = 10

Implicit weights are used for B-slices when WeightedBiPredFlag = 10. In this mode the offsets are always zero and the weights are 9-bits. To fit this in the same memory footprint, the offsets are not sent and the 9-bit weights are sign extended into the 16-bit block used for the weight/offset pair in explicit mode.

Indirect data Implicit Weight block in AVC-IT mode (and AVC-MC mode)

| DWord | Bit | MvSize | | |
|-------|-------|--------|---------------|---------------------|
| | | 0 | 2 | 8, 16, 32 |
| | 31:16 | n/a | Weight_Y_L1 | Weight_Y_Block0_L1 |
| | 15:0 | n/a | Weight_Y_L0 | Weight_Y_Block0_L0 |
| 1 | 31:16 | n/a | Weight_Cb_L1 | Weight_Cb_Block0_L1 |
| | 15:0 | n/a | Weight_Cb_L0 | Weight_Cb_Block0_L0 |
| 2 | 31:16 | n/a | Weight_Cr_L1 | Weight_Cr_Block0_L1 |
| | 15:0 | n/a | Weight_Cr_L0 | Weight_Cr_Block0_L0 |
| 3 | 31:0 | n/a | Reserved: MBZ | Reserved: MBZ |
| 4 | 31:16 | n/a | n/a | Weight_Y_Block1_L1 |
| | 15:0 | n/a | n/a | Weight_Y_Block1_L0 |
| 5 | 31:16 | n/a | n/a | Weight_Cb_Block1_L1 |
| | 15:0 | n/a | n/a | Weight_Cb_Block1_L0 |
| 6 | 31:16 | n/a | n/a | Weight_Cr_Block1_L1 |
| | 15:0 | n/a | n/a | Weight_Cr_Block1_L0 |
| 7 | 31:0 | n/a | n/a | Reserved: MBZ |
| 8 | 31:16 | n/a | n/a | Weight_Y_Block2_L1 |
| | 15:0 | n/a | n/a | Weight_Y_Block2_L0 |
| 9 | 31:16 | n/a | n/a | Weight_Cb_Block2_L1 |
| | 15:0 | n/a | n/a | Weight_Cb_Block2_L0 |
| 10 | 31:16 | n/a | n/a | Weight_Cr_Block2_L1 |
| | 15:0 | n/a | n/a | Weight_Cr_Block2_L0 |
| 11 | 31:0 | n/a | n/a | Reserved: MBZ |
| 12 | 31:16 | n/a | n/a | Weight_Y_Block3_L1 |
| | 15:0 | n/a | n/a | Weight_Y_Block3_L0 |
| 13 | 31:16 | n/a | n/a | Weight_Cb_Block3_L1 |
| | 15:0 | n/a | n/a | Weight_Cb_Block3_L0 |
| 14 | 31:16 | n/a | n/a | Weight_Cr_Block3_L1 |
| | 15:0 | n/a | n/a | Weight_Cr_Block3_L0 |



| DWord | Bit | MvSize | | |
|-------|------|--------|-----|---------------|
| | | 0 | 2 | 8, 16, 32 |
| 15 | 31:0 | n/a | n/a | Reserved: MBZ |

The weights for MvSize = 8,16,32 are replicated in exactly the same manner as the binding table indices. See section 1.7.3.1.5 for the description of the replication method. For MvSize=2 the replication is described in the following table:

| MbType | Inter Macroblock | Implicit Weight Replication Rule |
|--------|------------------|---|
| 1 | BP_L0_16x16 | L0 weight and offset replicated to L1 entries |
| 2 | B_L1_16x16 | L1 weight and offset replicated to L0 entries |
| 3 | B_Bi_16x16 | No replication needed. |

1.7.3.2.4 Transform Residual Block of Indirect Data in AVC-IT Mode

Transform-domain residual data block in AVC-IT mode is similar to that in IS mode. Only the non-zero coefficients are present in the data buffer and they are packed in the 8x8 block sequence of Y0, Y1, Y2, Y3, Cb4 and Cr5, as shown in Figure 1-4. When an 8x8 block is further subdivided into 4x4 subblocks, the coefficients, if present, are organized in the subblock order. The smallest subblock division is referred to as a **transform block**. The indirect data length in MEDIA_OBJECT_EX includes all the non-zero coefficients for the macroblock. It must be doubleword aligned.

Each non-zero coefficient in the indirect data buffer is contained in a doubleword-size data structure consisting of the coefficient index, end of block (EOB) flag and the fixed-point coefficient value in 2's complement form. As shown in **Error! Reference source not found.**, *index* is the row major 'raster' index of the coefficient within a transform block. A coefficient is a 16-bit value in 2's complement.

Structure of a transform-domain residue unit

| DWord | Bit | Description |
|-------|-------|--|
| 0 | 31:16 | Transform-Domain Residual (coefficient) Value. This field contains the value of the non-zero transform-domain residual data in 2's complement. |
| | 15:7 | Reserved: MBZ |
| | 6:1 | Index. This field specifies the raster-scan address (raw address) of the coefficient within the transform block. For a coefficient at Cartesian location (row, column) = (y, x) in a transform block of width W, Index is equal to (y * W + x). For example, coefficient at location (row, column) = (0, 0) in a 4x4 transform block has an index of 0; that at (2, 3) has an index of 2*4 + 3 = 11. The valid range of this field depends on the size of the transform block. Format = U6 Range = [0, 63] |
| | 0 | EOB (End of Block). This field indicates whether the transform-domain residue is the last one of the current transform block. |



1.7.3.3 Inline Data Format in AVC-MC Mode

Each MEDIA_OBJECT_EX command in “AVC-MC mode” corresponds to the processing of one macroblock. Macroblock parameters are passed in as inline data and the pixel-domain residual data (as well as motion vectors and weight/offset) for the macroblock is passed in as indirect data.

Inline data format in AVC-MC mode follows the exact same format like the one in AVC-IT mode. Specifically, the common fields required by VFE are at the same locations and have the same meaning.

Error! Reference source not found. depicts the inline data format in AVC-MC mode. Unlike AVC-IT, all fields in inline data are forwarded to the thread. Starting at GRF location, inline data are stored in GRF contiguously with the tail-end partial GRF, if present, zero-filled. Some fields are merely forwarded. Some fields are also used by VFE as indicated in the following table by a mark of [Used by VFE]. As shown, inline data starts at dword 4 of MEDIA_OBJECT_EX command.

Inline data in AVC-MC mode ([DevCTG])

| DWord | Bit | Description |
|-------|-------|---|
| 4+0 | 31:27 | Reserved. MBZ |
| | 26:25 | MbAffFieldFlag This field indicates that the current macroblock is a field macroblock within a MbAff frame picture. It is provided as Flag = MbaffFrame & FieldMbFlag . 00 = if (Flag == 0) 11 = if (Flag == 1) Other encodings are reserved |
| | 24 | FieldMbPolarityFlag This field indicates the field polarity of the current macroblock. Within a MbAff frame picture, this field may be different per macroblock and is set to 1 only for the second macroblock in a MbAff pair if FieldMbFlag is set. Otherwise, it is set to 0. Within a field picture, this field is set to 1 if the current picture is the bottom field picture. Otherwise, it is set to 0. It is a constant for the whole field picture. This field is reserved and MBZ for a progressive frame picture. 0 = Current macroblock is a field macroblock from the top field 1 = Current macroblock is a field macroblock from the bottom field <i>Programming Note: Here bits [26:24] (MbAffFieldFlag and FiedMbPolarityFlag) match with bits [10:8] of the Media Block Read message descriptor, simplifying the programming for message generation, as when MbAffFieldFlag is “1”, kernels need to override the original “frame” surface state set for MBAFF frame picture.</i> |
| | 23 | Reserved: MBZ |



| DWord | Bit | Description |
|-------|-------|--|
| | 22:20 | <p>MvSize (Motion Vector Size) [Used by VFE]. This field specifies the size of motion vectors for the macroblock stored in the indirect data buffer. The valid numbers are listed below indicating the size of the regrouped motion vectors.</p> <p>This field is reserved (MBZ) when IntraMbFlag = 1.</p> <p>000 = 0: No motion vector 001 reserved 010 = 2MV: One motion vector pair 011 reserved 100 = 8MV: Four motion vector pairs 101 = 16MV: 16 motion vectors 110 = 32MV: 16 motion vector pairs 111 reserved</p> <p><i>This field is identical to that in AVC-IT mode.</i></p> |
| | 19:16 | Reserved. MBZ |
| | 15 | <p>Transform8x8Flag. This field equals to the value of <i>transform_size_8x8_flag</i> as defined in AVC spec. If set, it indicates that luma samples are in residual 8x8 blocks. Otherwise, it indicates that luma samples are in residual 4x4 blocks.</p> <p>0: luma residual 4x4 blocks 1: luma residual 8x8 blocks</p> |
| | 14 | <p>FieldMbFlag (Field Macroblock Flag). This field specifies whether current macroblock is field macroblock.</p> <p>0 = Frame macroblock. 1 = Field macroblock.</p> |
| | 13 | <p>IntraMbFlag (Intra Macroblock Flag). This field specifies whether the current macroblock is an Intra (I) macroblock. A collective macroblock type in AVC standard includes I, SI, P and B. Type SI is not supported.</p> <p>0 = P or B macroblock. 1 = I macroblock.</p> |
| | 12:8 | <p>MbType (Macroblock Type). This field, along with IntraMbFlag specifies the macroblock types.</p> <p>Further details can be found in Section 1.7.3.1.2.</p> |
| | 7:6 | <p>WeightedBiPredFlag (Weighted Bidirectional Prediction Flag) (from Picture State). Valid only for macroblock in inter mode. Otherwise (intra macroblock), this field is reserved.</p> |
| | 5 | <p>WeightedPredFlag (Weighted Prediction Flag) [from Picture State]. Valid only for macroblock in inter mode. Otherwise (intra macroblock), this field is reserved.</p> |
| | 4 | Reserved. MBZ |
| | 3:2 | <p>ChromaFormatIdc (Chroma Format Indicator). This field is equal to the value of <i>chroma_format_idct</i> as defined in AVC (§7.4.2.1). It specifies the chroma sampling relative to the luma sampling.</p> <p>This field is constant within a picture.</p> <p>00 = Luma only (monochrome) 01 = YUV420 sampling 10 is reserved (for YUV422 sampling) 11 is reserved (for YUV444 sampling)</p> |



| DWord | Bit | Description |
|-------|-------|--|
| | 1 | <p>MbaffFrameFlag (MB-AFF Frame Flag) [PPS]. This field indicates whether the current picture is a progressive frame or a MB-AFF (macroblock adaptive frame field) frame picture.</p> <p>This field is frame Picture Parameter Set.</p> <p>This field is reserved (MBZ) if FieldPicFlag = 1.</p> <p>0 = Progressive frame picture 1 = MB-AFF frame picture</p> |
| | 0 | <p>FieldPicFlag (Field Picture Flag) [PPS]. This field indicates whether the current picture is a field or a frame picture. A frame picture may be a progressive frame picture or an MB-AFF frame picture depending on the value of MbaffFrame.</p> <p>This field is frame Picture Parameter Set.</p> <p>0 = Frame picture 1 = Field picture</p> |
| 4+1 | 31:16 | <p>CbpY (Coded Block Pattern Y)</p> <p><i>Not expected to be used by Kernel. Thus, programming this field by host software is optional.</i></p> |
| | 15:8 | <p>VertOrigin (Vertical Origin). This field specifies the vertical origin of current macroblock in the destination picture in units of macroblocks. For field macroblock pair in MBAFF frame, the vertical origins for both macroblocks should be set as if they were located in corresponding field pictures. For example, for field macroblock pair originated at (16, 64) pixel location in an MBAFF frame picture, the Vertical Origin for both macroblocks should be set as 2 (macroblocks).</p> <p>Format = U8 in unit of macroblock.</p> |
| | 7:0 | <p>HorzOrigin (Horizontal Origin). This field specifies the horizontal origin of current macroblock in the destination picture in units of macroblocks.</p> <p>Format = U8 in unit of macroblock.</p> |
| 4+2 | 31:14 | <p>Reserved. MBZ</p> |
| | 13:8 | <p>Coded Block Pattern [Used by VFE). This field specifies whether blocks (8x8) are present or not.</p> <p>Each bit corresponds to one block. "0" indicates error block isn't present, "1" indicates error block is present.</p> <p>Bit 13: Y0 Bit 12: Y1 Bit 11: Y2 Bit 10: Y3 Bit 9: Cb Bit 8: Cr</p> <p><i>This field is different than the input in AVC-IT mode, but it is the same as the corresponding Root Thread payload field in AVC-IT mode.</i></p> |
| | 7:4 | <p>CbpCr</p> <p><i>Not expected to be used by Kernel. Thus, programming this field by host software is optional.</i></p> |
| | 3:0 | <p>CbpCb</p> <p><i>Not expected to be used by Kernel. Thus, programming this field by host software is optional.</i></p> |



| DWord | Bit | Description |
|------------------|------|--|
| 4+3 to 4+5 | 31:0 | For intra macroblocks, see Error! Reference source not found. in Section 0(Inline Data Format in AVC-IT Mode) For intra macroblocks, see in Section 0 (Inline Data Format in AVC-IT Mode) <i>This field is identical to that in AVC-IT mode.</i> |

Inline data in AVC-MC mode ([DevILK])

| DWord | Bit | Description |
|-------|----------------------|--|
| 4+0 | 31:27 | Reserved. MBZ |
| | 26:25 | MbAffFieldFlag This field indicates that the current macroblock is a field macroblock within a MbAff frame picture. It is provided as Flag = MbaffFrame & FieldMbFlag . 00 = if (Flag == 0) 11 = if (Flag == 1) Other encodings are reserved |
| | 24 | FieldMbPolarityFlag This field indicates the field polarity of the current macroblock. Within a MbAff frame picture, this field may be different per macroblock and is set to 1 only for the second macroblock in a MbAff pair if FieldMbFlag is set. Otherwise, it is set to 0. Within a field picture, this field is set to 1 if the current picture is the bottom field picture. Otherwise, it is set to 0. It is a constant for the whole field picture. This field is reserved and MBZ for a progressive frame picture. 0 = Current macroblock is a field macroblock from the top field 1 = Current macroblock is a field macroblock from the bottom field <i>Programming Note: Here bits [26:24] (MbAffFieldFlag and FiedlMbPolarityFlag) match with bits [10:8] of the Media Block Read message descriptor, simplifying the programming for message generation, as when MbAffFieldFlag is "1", kernels need to override the original "frame" surface state set for MBAFF frame picture.</i> |
| 23 | Reserved: MBZ | |



| DWord | Bit | Description |
|-------|-------|---|
| | 22:20 | <p>MvSize (Motion Vector Size) [Used by VFE]. This field specifies the size of motion vectors for the macroblock stored in the indirect data buffer. The valid numbers are listed below indicating the size of the regrouped motion vectors. Details are provided in Section 1.7.3.1.4.</p> <p>This field is reserved (MBZ) when IntraMbFlag = 1.</p> <p>000 = 0: No motion vector 001 reserved 010 = 2MV: One motion vector pair 011 reserved 100 = 8MV: Four motion vector pairs 101 = 16MV: 16 motion vectors 110 = 32MV: 16 motion vector pairs 111 reserved</p> |
| | 19 | DcBlockCodedYFlag [Used by VFE]. This field is consumed by VFE and is not delivered to the thread. |
| | 18 | DcBlockCodedCbFlag [Used by VFE]. This field is consumed by VFE and is not delivered to the thread. |
| | 17 | DcBlockCodedCrFlag [Used by VFE]. This field is consumed by VFE and is not delivered to the thread. |
| | 16 | Reserved. MBZ |
| | 15 | <p>Transform8x8Flag [Used by VFE]. This field equals to the value of <i>transform_size_8x8_flag</i> as defined in AVC spec. If set, it indicates that luma samples are in residual 8x8 blocks. Otherwise, it indicates that luma samples are in residual 4x4 blocks.</p> <p>0: luma residual 4x4 blocks 1: luma residual 8x8 blocks</p> |
| | 14 | <p>FieldMbFlag (Field Macroblock Flag). This field specifies whether current macroblock is field macroblock.</p> <p>0 = Frame macroblock. 1 = Field macroblock.</p> |
| | 13 | <p>IntraMbFlag (Intra Macroblock Flag). This field specifies whether the current macroblock is an Intra (I) macroblock. A collective macroblock type in AVC standard includes I, SI, P and B. Type SI is not supported.</p> <p>0 = P or B macroblock. 1 = I macroblock.</p> |
| | 12:8 | <p>MbType (Macroblock Type). This field, along with IntraMbFlag specifies the macroblock types.</p> <p>Further details can be found in Section 1.7.3.1.2.</p> |



| DWord | Bit | Description |
|-------|-------|---|
| | 7:6 | <p>WeightedBiPredFlag (Weighted Bidirectional Prediction Flag) (from Picture State). This field specifies the bidirectional prediction mode and is derived from syntax elements <i>weighted_bipred_flag</i> and <i>weighted_pred_flag</i> as defined in AVC spec.</p> <p>It is valid only for inter predicted macroblock. Otherwise (intra macroblock), this field is reserved and MBZ.</p> <p>For B-macroblock, this field is the same as <i>weighted_bipred_flag</i> as defined in AVC spec.</p> <p>00 = Default weighted prediction 01 = Explicit weighted prediction 10 = Implicit weighted prediction 11 = Reserved.</p> <p>For P-macroblock, the MSB is always 0 and the LSB is the same as <i>weighted_pred_flag</i> as defined in AVC spec.</p> <p>00 = Default weighted prediction 01 = Explicit weighted prediction 10 and 11 are reserved</p> |
| | 5 | <p>WeightedPredFlag (Weighted Prediction Flag) [from Picture State]. It is valid only for inter predicted macroblock. Otherwise (intra macroblock), this field is reserved and MBZ.</p> <p>0 = Default weighted prediction 1 = Explicit weighted prediction</p> <p><i>Note: Information in this field is also carried in WeightedBiPredFlag.</i></p> |
| | 4 | Reserved. MBZ |
| | 3:2 | <p>ChromaFormatIdc (Chroma Format Indicator). This field is equal to the value of <i>chroma_format_idct</i> as defined in AVC (§7.4.2.1). It specifies the chroma sampling relative to the luma sampling.</p> <p>This field is constant within a picture.</p> <p>00 = Luma only (monochrome) 01 = YUV420 sampling 10 is reserved (for YUV422 sampling) 11 is reserved (for YUV444 sampling)</p> |
| | 1 | <p>MbaffFrameFlag (MB-AFF Frame Flag) [PPS]. This field indicates whether the current picture is a progressive frame or a MB-AFF (macroblock adaptive frame field) frame picture.</p> <p>This field is frame Picture Parameter Set.</p> <p>This field is reserved (MBZ) if FieldPicFlag = 1.</p> <p>0 = Progressive frame picture 1 = MB-AFF frame picture</p> |
| | 0 | <p>FieldPicFlag (Field Picture Flag) [PPS]. This field indicates whether the current picture is a field or a frame picture. A frame picture may be a progressive frame picture or an MB-AFF frame picture depending on the value of MbaffFrame.</p> <p>This field is frame Picture Parameter Set.</p> <p>0 = Frame picture 1 = Field picture</p> |
| 4+1 | 31:16 | CbpY (Coded Block Pattern Y) [Used by VFE] |



| DWord | Bit | Description |
|------------------|-------|---|
| | 15:8 | VertOrigin (Vertical Origin). This field specifies the vertical origin of current macroblock in the destination picture in units of macroblocks. For field macroblock pair in MBAFF frame, the vertical origins for both macroblocks should be set as if they were located in corresponding field pictures. For example, for field macroblock pair originated at (16, 64) pixel location in an MBAFF frame picture, the Vertical Origin for both macroblocks should be set as 2 (macroblocks). Format = U8 in unit of macroblock. |
| | 7:0 | HorzOrigin (Horizontal Origin). This field specifies the horizontal origin of current macroblock in the destination picture in units of macroblocks. Format = U8 in unit of macroblock. |
| 4+2 | 31:24 | QpPrimeCr [Used by VFE] |
| | 23:16 | QpPrimeCb [Used by VFE] |
| | 15:8 | QpPrimeY [Used by VFE] |
| | 7:4 | CbpCr [Used by VFE] |
| | 3:0 | CbpCb [Used by VFE] |
| 4+3 to 4+5 | 31:0 | For intra macroblocks, see Error! Reference source not found. For intra macroblocks, see |
| 4+6 | 31:16 | LevelScaleCb. [Used by VFE] This field is for inverse transform of the Chroma (Cb) DC block. The LevelScale field is consumed by VFE and is not needed by the thread, but since the GRF has to be filled to the end of the block it should be sent anyways. |
| | 15:0 | LevelScaleCr. [Used by VFE] This field is for inverse transform of the Chroma (Cr) DC block. |
| 4+7 | 31:16 | Reserved. MBZ |
| | 15:0 | LevelScaleY. [Used by VFE] This field is for inverse transform of the Luma DC block. |
| 4+8 | 31:0 | BSD_Pass_Thru_Data |
| 4+9 | 31:0 | BSD_Pass_Thru_Data |
| 4+10 | 31:0 | BSD_Pass_Thru_Data |
| 4+11 | 31:0 | BSD_Pass_Thru_Data |

1.7.3.4 Indirect Data Format in AVC-MC Mode

Indirect data in AVC-IT mode consist of Motion Vectors, Weight/Offset and pixel-domain residual data. All three data blocks have variable size.



Sizes of Motion Vector block and the Weight-Offset block are determined by the MvSize value. They are the same as in AVC-IT mode and are depicted in **Error! Reference source not found.** Weight-Offset block, if present, is always packed behind the Motion Vector block. See Section 1.7.3.2 for more details.

Residual data block must start at a predetermined offset, controlled by the fields in VFE_STATE_EX.

1.7.3.5 Inline Data Format in VC1-IT Mode

Each MEDIA_OBJECT_EX command in “VC1-IT mode” corresponds to the processing of one macroblock. Macroblock parameters, including motion vectors, are passed in as inline data and the non-zero DCT coefficient data for the macroblock is passed in as indirect data.

Error! Reference source not found. depicts the inline data format in VC1-IT mode. All fields in inline data are forwarded to the thread as thread payload. Inline data are stored in GRF contiguously with the tail-end partial GRF, if present, zero-filled. Some fields are merely forwarded. Some fields are also used by VFE as indicated in the following table by a mark of [Used by VFE]. As shown, inline data starts at dword 4 of MEDIA_OBJECT_EX command.

Inline data in VC1-IT mode

| DWord | Bit | Description | | | | | | | | | | |
|-------|--|---|-----|-------------|----|---|----|--|----|---|----|--|
| 4+0 | 31:28 | <p>MvFieldSelect. A bit-wise representation indicating which field in the reference frame is used as the reference field for current field. It's only used in decoding interlaced pictures. This field is valid for non-intra macroblock only.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>28</td> <td>Forward predict of current frame/field or TOP field of interlace frame, or block 0 in 4MV mode.</td> </tr> <tr> <td>29</td> <td>Backward predict of current frame/field or TOP field of interlace frame, or forward predict for block 1 in 4MV mode.</td> </tr> <tr> <td>30</td> <td>Forward predict of BOTTOM field of interlace frame, or block 2 in 4MV mode.</td> </tr> <tr> <td>31</td> <td>Backward predict of BOTTOM field of interlace frame, or forward predict for block 3 in 4MV mode.</td> </tr> </tbody> </table> <p>Each corresponding bit has the following indication. 0 = The prediction is taken from the <u>top</u> reference field. 1 = The prediction is taken from the <u>bottom</u> reference field.</p> | Bit | Description | 28 | Forward predict of current frame/field or TOP field of interlace frame, or block 0 in 4MV mode. | 29 | Backward predict of current frame/field or TOP field of interlace frame, or forward predict for block 1 in 4MV mode. | 30 | Forward predict of BOTTOM field of interlace frame, or block 2 in 4MV mode. | 31 | Backward predict of BOTTOM field of interlace frame, or forward predict for block 3 in 4MV mode. |
| Bit | Description | | | | | | | | | | | |
| 28 | Forward predict of current frame/field or TOP field of interlace frame, or block 0 in 4MV mode. | | | | | | | | | | | |
| 29 | Backward predict of current frame/field or TOP field of interlace frame, or forward predict for block 1 in 4MV mode. | | | | | | | | | | | |
| 30 | Forward predict of BOTTOM field of interlace frame, or block 2 in 4MV mode. | | | | | | | | | | | |
| 31 | Backward predict of BOTTOM field of interlace frame, or forward predict for block 3 in 4MV mode. | | | | | | | | | | | |
| | 27 | Reserved. MBZ | | | | | | | | | | |
| | 26 | <p>MvFieldSelectChroma . This field specifies the polarity of reference field for chroma blocks when their motion vector is derived in Motion4MV mode for interlaced (field) picture. Non-intra macroblock only. This field is derived from MvFieldSelect. 0 = The prediction is taken from the <u>top</u> reference field. 1 = The prediction is taken from the <u>bottom</u> reference field.</p> | | | | | | | | | | |



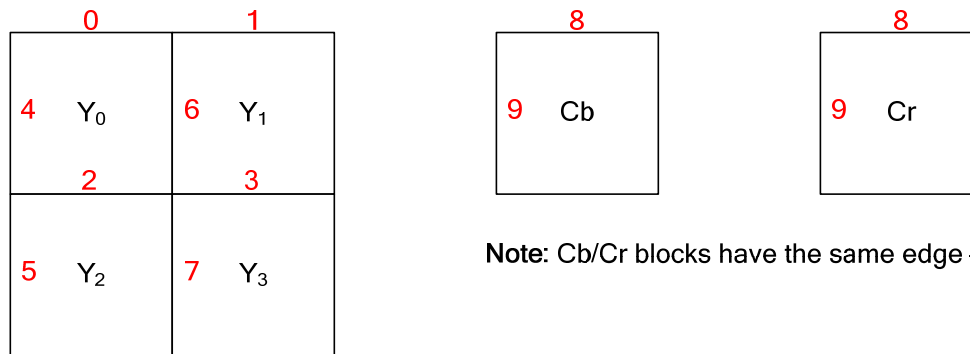
| DWord | Bit | Description |
|-------|-------|--|
| | 25:24 | <p>MotionType – Motion Type</p> <p>For frame picture, a macroblock may only be either 00 or 10.</p> <p>For interlace picture, a macroblock may be of any motion types. It can be 01 if and only if DctType is 1.</p> <p>This field is 00 if and only if IntraMacroblock is 1.</p> <p>00 = Intra 01 = Field Motion. 10 = Frame Motion or no motion. Others = Reserved.</p> |
| | 23 | Reserved. MBZ |
| | 22 | <p>MvSwitch. This field specifies whether the prediction needs to be switched from forward to backward or vice versa for single directional prediction for top and bottom fields of interlace frame B macroblocks.</p> <p>0 = No directional prediction switch from top field to bottom field 1 = Switch directional prediction from top field to bottom field</p> |
| | 21 | <p>DctType. This field specifies whether the residual data is coded as field residual or frame residual for interlaced picture. This field can be 1 only if MotionType is 00 (intra) or 01 (field motion).</p> <p>For progressive picture, this field must be set to '0', i.e. all macroblocks are frame macroblock.</p> <p>0 = Frame residual type. 1 = Field residual type.</p> |
| | 20 | <p>OverlapTransform. This field indicates whether overlap smoothing filter should be performed on I-block boundaries.</p> <p>0 = No overlap smoothing filter. 1 = Overlap smoothing filter performed.</p> |
| | 19 | <p>Motion4MV. This field indicates whether current macroblock a progressive P picture uses 4 motion vectors, one for each luminance block.</p> <p>It's only valid for progressive P-picture decoding. Otherwise, it is reserved and MBZ. For example, with MotionForward is 0, this field must also be set to 0.</p> <p>0 = 1MV-mode. 1 = 4MV-mode.</p> |
| | 18 | <p>MotionBackward. This field specifies whether the backward motion vector is active for B-picture. This field must be 0 if Motion4MV is 1 (no backward motion vector in 4MV-mode).</p> <p>0 = No backward motion vector. 1 = Use backward motion vector(s).</p> |
| | 17 | <p>MotionForward. This field specifies whether the forward motion vector is active for P and B pictures.</p> <p>0 = No forward motion vector. 1 = Use forward motion vector(s).</p> |



| DWord | Bit | Description |
|-------|-------|--|
| | 16 | <p>IntraMacroblock. This field specifies if the current macroblock is intra-coded. When set, Coded Block Pattern is ignored and no prediction is performed (i.e., no motion vectors are used).</p> <p>For field motion, this field indicates whether the top field of the macroblock is coded as intra.</p> <p>0 = Non-intra macroblock. 1 = Intra macroblock.</p> |
| | 15:12 | <p>LumaIntra8x8Flag – Luma Intra 8x8 Flag</p> <p>This field specifies whether each of the four 8x8 luminance blocks are intra or inter coded when Motion4MV is set to 4MV-Mode.</p> <p>Each bit corresponds to one block. “0” indicates the block is inter coded and ‘1’ indicates the block is intra coded.</p> <p>When Motion4MV is not 4MV-Mode, this field is reserved and MBZ.</p> <p>Bit 15: Y0 Bit 14: Y1 Bit 13: Y2 Bit 12: Y3</p> |
| | 11:6 | <p>CBP - Coded Block Pattern</p> <p>This field specifies whether the 8x8 residue blocks in the macroblock are present or not. Each bit corresponds to one block. “0” indicates residue block isn’t present, “1” indicates residue block is present.</p> <p>Note: For each block in an intra-coded macroblock or an intra-coded block in a P macroblock in 4MV-Mode, the corresponding CBP must be 1. Subsequently, there must be at least one coefficient (this coefficient might be zero) in the indirect data buffer associated with the block (i.e. residue block must be present).</p> <p>Bit 11: Y0 Bit 10: Y1 Bit 9: Y2 Bit 8: Y3 Bit 7: Cb4 Bit 6: Cr5</p> |
| | 5 | <p>ChromaIntraFlag - Derived Chroma Intra Flag</p> <p>This field specifies whether the chroma blocks should be treated as intra blocks based on motion vector derivation process in 4MV mode.</p> <p>0 = Chroma blocks are not coded as intra. 1 = Chroma blocks are coded as intra</p> |
| | 4 | <p>LastRowFlag – Last Row Flag</p> <p>This field indicates that the current macroblock belongs to the last row of the picture. This field may be used by the kernel to manage pixel output when overlap transform is on.</p> <p>0 = Not in the last row 1 = In the last row</p> |
| | 3:0 | Reserved. MBZ |
| 4+1 | 32:26 | Reserved. MBZ |



| DWord | Bit | Description | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|-------|---|--|---------------------------------------|--|--|------------|------------|--|------|---------|-------|-------|-------|-------|----|------------------------|-----|------------|------------|------------|----|---------------------------|-----|-----|------------|------------|----|---------------------------|-----|-----|------------|------------|----|----------------------------|-----|-----|-----|-----|
| | 15:8 | VertOrigin (Vertical Origin) In unit of macroblocks relative to the current picture (frame or field). | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 7:0 | HorzOrigin (Horizontal Origin) In unit of macroblocks. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4+2 | 31:16 | MotionVector[0].Vert | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 15:0 | MotionVector[0].Horz | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4+3 | 31:0 | MotionVector[1] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4+4 | 31:0 | MotionVector[2] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4+5 | 31:0 | MotionVector[3] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4+6 | 31:0 | MotionVectorChroma OPEN: This field is not valid for a field motion in an interlaced frame picture where 4 MVs for chroma blocks. Notes: This field is derived from MotionVector[3:0] as described in the following section. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4+7 | 32:24 | Subblock Code for Y3 [Used by VFE] The following subblock coding definition applies to all 6 subblock coding bytes. Bits 7:6 are reserved. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | <table border="1"> <thead> <tr> <th colspan="2">Subblock Partitioning (Bits [1:0])</th> <th colspan="4">Subblock Present (0 means not present, 1 means present)</th> </tr> <tr> <th>Code</th> <th>Meaning</th> <th>Bit 2</th> <th>Bit 3</th> <th>Bit 4</th> <th>Bit 5</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Single 8x8 block (sb0)</td> <td>Sb0</td> <td>Don't care</td> <td>Don't care</td> <td>Don't care</td> </tr> <tr> <td>01</td> <td>Two 8x4 subblocks (sb0-1)</td> <td>Sb0</td> <td>Sb1</td> <td>Don't care</td> <td>Don't care</td> </tr> <tr> <td>10</td> <td>Two 4x8 subblocks (sb0-1)</td> <td>Sb0</td> <td>Sb1</td> <td>Don't care</td> <td>Don't care</td> </tr> <tr> <td>11</td> <td>Four 4x4 subblocks (sb0-3)</td> <td>Sb0</td> <td>Sb1</td> <td>Sb2</td> <td>Sb3</td> </tr> </tbody> </table> | Subblock Partitioning (Bits [1:0]) | | Subblock Present (0 means not present, 1 means present) | | | | Code | Meaning | Bit 2 | Bit 3 | Bit 4 | Bit 5 | 00 | Single 8x8 block (sb0) | Sb0 | Don't care | Don't care | Don't care | 01 | Two 8x4 subblocks (sb0-1) | Sb0 | Sb1 | Don't care | Don't care | 10 | Two 4x8 subblocks (sb0-1) | Sb0 | Sb1 | Don't care | Don't care | 11 | Four 4x4 subblocks (sb0-3) | Sb0 | Sb1 | Sb2 | Sb3 |
| | | Subblock Partitioning (Bits [1:0]) | | Subblock Present (0 means not present, 1 means present) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | Code | Meaning | Bit 2 | Bit 3 | Bit 4 | Bit 5 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 00 | Single 8x8 block (sb0) | Sb0 | Don't care | Don't care | Don't care | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 01 | Two 8x4 subblocks (sb0-1) | Sb0 | Sb1 | Don't care | Don't care | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 10 | Two 4x8 subblocks (sb0-1) | Sb0 | Sb1 | Don't care | Don't care | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 11 | Four 4x4 subblocks (sb0-3) | Sb0 | Sb1 | Sb2 | Sb3 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 23:16 | Subblock Code for Y2 [Used by VFE] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 15:8 | Subblock Code for Y1 [Used by VFE] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 7:0 | Subblock Code for Y0 [Used by VFE] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4+8 | 31:16 | Reserved. MBZ | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 15:8 | Subblock Code for Cr [Used by VFE] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 7:0 | Subblock Code for Cb [Used by VFE] | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |



Indexing Block Edges for Overlapped Smoothing

1.7.3.5.1 Deriving Motion Vectors and Field Select for Interlaced Frame Picture

In MPEG2, the motion vectors are related to the decoded picture. For field picture, it is related to the field and which field of the reference frame that a motion vector points to is given in the bitstream by syntax element called Motion Vector Field Select (MVFS). In contrary, motion vectors defined in VC1 standard for an interlaced frame is frame based and there is no such MVFS syntax. The VC1-IT interface defines the motion vector and MVFS following the MPEG2 convention. Therefore, motion vectors and MVFS must be derived from the frame-based motion vector values and the current macroblock position (in the top or bottom field).

The derivation of the picture-based motion vectors (luma) and MVFS is provided by the following pseudo-code. The idea is that MVFS comes from the LSB of the final pointer to the reference frame (note here it is frame based not field base). The final pointer is the addition of the frame based motion vector and the current macroblock position in the current frame (again, it is relative to the frame, not picture).

- Let (MV_X, MV_Y) be the original frame based luma motion vector in quarter-pel representation.
- Let (LMV_X, LMV_Y) be the derived field based luma motion vector and (CMV_X, CMV_Y) be the derived field based chroma motion vector, both in quarter-pel precision as well.
- Let MVFS be the derived motion vertical field select field.
- Then
 - LMV_X = MV_X;
 - if (Current_field != BOTTOM_FIELD)
 - iy = MV_Y >> 2; // Interger portion of MV_Y
 - else // Current_field == BOTTOM_FIELD
 - iy = (MV_Y >> 2) + 1; // Interger portion of MV_Y adjusted
 - MVFS = iy & 1; // 0 – top field, 1 – bottom field
 - LMV_Y = ((iy >> 1) << 2) + (MV_Y & 3);



1.7.3.5.2

1.7.3.5.3 Chroma Interpolations for Motion Prediction

There are two different interpolation modes are used for generating chroma samples according to the picture level parameter **BMVprecisionAndChromaRelation**: *the quarter-pel chroma motion prediction*, and *the half-pel chroma motion prediction*. The bilinear interpolation is applied for both cases.

For the quarter-pel case, the motion vectors for the chroma components are derived from the corresponding luma motion vectors according to the following pseudocodes.

For the case of 1-MV,

```
cmv.x = (mv.x + (mv.x&3==3))>>1;
cmv.y = (mv.y + (mv.y&3==3))>>1;
```

For the case of 4-MV,

```
switch(number of inter-coded blocks)
case 3: // median of three
    if(mv0.x<mv1.x && mv0.x<mv2.x)
        mv0.x = (mv1.x<mv2.x) ? mv1.x : mv2.x;
    else if(mv0.x>mv1.x && mv0.x>mv2.x)
        mv0.x = (mv1.x>mv2.x) ? mv1.x : mv2.x;
    cmv.x = (mv0.x + (mv0.x&3==3))>>1;
    cmv.y = (mv0.y + (mv0.y&3==3))>>1;
    break;
case 4: // average of middle two
    if(mv0.x<mv1.x && mv0.x<mv2.x && mv0.x<mv3.x){
        if(mv2.x<mv3.x){ mv0.x = mv2.x; if(mv1.x>mv3.x) mv1.x=mv3.x; }
        else { mv0.x = mv3.x; if(mv1.x>mv2.x) mv1.x=mv2.x; }
    }
    else if(mv0.x>mv1.x && mv0.x>mv2.x && mv0.x>mv3.x){
        if(mv2.x>mv3.x){ mv0.x = mv2.x; if(mv1.x<mv3.x) mv1.x=mv3.x; }
        else { mv0.x = mv3.x; if(mv1.x<mv2.x) mv1.x=mv2.x; }
    }
    else if(mv1.x<mv2.x && mv1.x<mv3.x)
        mv1.x = (mv2.x<mv3.x) ? mv2.x : mv3.x;
    else if(mv1.x>mv2.x && mv1.x>mv3.x)
        mv1.x = (mv2.x>mv3.x) ? mv2.x : mv3.x;
case 2: // average of two
    x = (mv0.x + mv1.x)>>1; cmv.x = (x + (x&3==3))>>1;
    y = (mv0.y + mv1.y)>>1; cmv.y = (y + (y&3==3))>>1;
    break;
case 0: case 1: chroma should be coded as intra-blocks.
}
```



For the half-pel case, the motion vectors for the chroma components are derived from one more extra shifting operation by rounding to the nearest full-pel if they are not currently in the half-grid.

For simple and main profile, the motion vectors are truncated so that the reference block is not totally off the picture frame:

```
// For luma:
if((mv.x>>2)<-16) mv.x = -64      +(mv.x&3);
if((mv.x>>2)> PW) mv.x = (PW<<2)+(mv.x&3);
if((mv.y>>2)<-16) mv.y = -64      +(mv.y&3);
if((mv.y>>2)> PH) mv.y = (PH<<2)+(mv.y&3);

// For chroma:
if((cmv.x>>2)<- 8) cmv.x = -32      +(cmv.x&3);
if((cmv.x>>2)>CPW) cmv.x = (CPW<<2)+(cmv.x&3);
if((cmv.y>>2)<- 8) cmv.y = -32      +(cmv.y&3);
if((cmv.y>>2)>CPH) cmv.y = (CPH<<2)+(cmv.y&3);
```

1.7.3.6 Indirect Data Format in VC1-IT Mode

Indirect data format in VC1-IT mode is identical to the transform-domain residual data block portion of the indirect data format in AVC-IT mode.

The indirect data start address in MEDIA_OBJECT_EX specifies the doubleword aligned address of the first non-zero transform-domain residue (referred to as ‘coefficient’) of the first block of the macroblock. The indirect data length in MEDIA_OBJECT_EX includes all the non-zero coefficients for the macroblock. It must be doubleword aligned.

Each non-zero coefficient in the indirect data buffer is contained in a doubleword-size data structure as shown in **Error! Reference source not found.**

1.7.3.7 Inline Data Format in Generic Mode [DebCTG+]

MEDIA_OBJECT_EX command can also be used in “Generic mode” in place of MEDIA_OBJECT command. The only difference of the usage is to allow interface descriptor remap. MEDIA_OBJECT_EX command cannot be used together with MEDIA_OBJECT command.

1.7.4 MEDIA_OBJECT_PRT Command [DevCTG+]

[DevBW/DevCL] This command is not supported.

The MEDIA_OBJECT_PRT command is for generating Persistent Root Thread for the media pipeline. It only supports loading of inline data but not indirect data.

This command should be used for a root thread that might have to be present in the system for a period longer than the certain minimal context-switch interrupt latency. It has to honor the context interrupt signal to terminate upon request. It should also handle replay from the interrupted point upon context restore (the same thread being dispatched more than once). In contrary, if a thread is not a Persistent Root Thread, if dispatched, it must run to completion.



The command can be used in all VFE modes, except VLD mode.

Note for [DevCTG+]: This command is supported for generating a general-purpose root thread independent of the VFE mode.

Note for [DevILK+]: Thread generated by this command always bypasses the hardware scoreboard if the hardware scoreboard is enabled.

[Pre-DevSNB]

| Dword | Bits | Description |
|-------|-------|---|
| 0 | 31:29 | Command Type = GFXPIPE = 3h |
| | 28:16 | Media Command Opcode = MEDIA_OBJECT_PRT Pipeline[28:27] = Media = 2h; Opcode[26:24] = 1h; Subopcode[23:16] = 02h |
| | 15:0 | DWord Length (Excludes DWords 0,1) Valid range: [3, 14] Note: Regardless of the mode, inline data must be present in this command. The command size must fit within 16 dwords. |
| 1 | 31:0 | Reserved. MBZ |
| 2 | 31 | Reserved. MBZ |
| | 30:24 | Interface Descriptor Offset. This field specifies the offset from the interface descriptor base pointer to the interface descriptor which will be applied to this object. It is specified in units of interface descriptors. Format = U7 |
| | 23 | PRT_Fence Needed. This field specifies that a PRT_Fence is generated after dispatching the thread associated with this MEDIA_OBJECT_PRT. The PRT_Fence prevents additional threads following this persistent root thread until a thread spawn message is sent. The PRT_Fence is generated on first dispatch of the persistent root, as well as on re-dispatches of the persistent root after context restore. Format = Enable |
| | 22:21 | Reserved. MBZ |
| | 20 | Thread Synchronization. This field when set indicates that the dispatch of the thread originated from this command is based on the "spawn root thread" message. <i>In VLD mode, this field must be programmed as 0, because the Children Present field in VFE_STATE must be 0 in this mode.</i> 0 = No thread synchronization 1 = Thread dispatch is synchronized by the "spawn root thread" message |
| | 19:0 | Reserved. MBZ (was Indirect Data Length) |
| 3 | 31:0 | Reserved. MBZ (was Indirect Data Start Address) |
| 4..N | 31:0 | Inline Data |



1.7.5 GPGPU_WALKER Command

| GPGPU_WALKER | | |
|---------------------|-------|--|
| Project: All | | Length Bias: 2 |
| Desc | | |
| DWord | Bit | Description |
| 0 | 31:29 | Command Type Default Value: 3h GFXPIPE Format: OpCode |
| | 28:27 | Pipeline Default Value: 2h Media Format: OpCode |
| | 26:24 | Media Command Opcode Default Value: 1h Format: OpCode |
| | 23:16 | SubOpcode A Default Value: 05h Format: OpCode |
| | 15:11 | Reserved Project: All Format: MBZ |
| | 10 | Indirect Parameter Enable Project: All Format: U1 If set, the values in DW 4, 6, 8 are ignored and replaced by the current values of the corresponding GPGPU_xxx MMIO registers: <ul style="list-style-type: none"> • GPGPU_DISPATCHDIMX (instead of DW4) • GPGPU_DISPATCHDIMY (instead of DW6) GPGPU_DISPATCHDIMZ (instead of DW8) |
| | 9 | Reserved Project: All Format: MBZ |
| | 8 | Predicate Enable Project: All Format: U1 If set, this command is executed (or not) depending on the current value of the MI Predicate internal state bit. This command is ignored only if PredicateEnable is set and the Predicate state bit is 0. |
| | 7:0 | DWord Length Default Value: 9h Allowed value is 9 Format: =n Total Length - 2 Project: All |
| 1 | 7:5 | Reserved Project: All Format: MBZ |
| | 4:0 | Interface Descriptor Offset Format: U5 FormatDesc This field specifies the offset from the interface descriptor base pointer to the interface descriptor which will be applied to this object. It is specified in units of interface descriptors. |



GPGPU_WALKER

| 2 | 31:30 | <p>SIMD Size</p> <p>This field determines the size of the payload and the number of bits of the execution mask that are expected. The kernel pointed to by the interface descriptor should match the SIMD declared here.</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">Value</th> <th style="text-align: center;">Name</th> <th style="text-align: center;">Description</th> <th style="text-align: center;">Project</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">0</td> <td style="text-align: center;">SIMD8</td> <td>8 LSBs of the execution mask are used</td> <td></td> </tr> <tr> <td style="text-align: center;">1</td> <td style="text-align: center;">SIMD16</td> <td>16 LSBs used in execution mask</td> <td></td> </tr> <tr> <td style="text-align: center;">2</td> <td style="text-align: center;">SIMD32</td> <td>32 bits of execution mask used</td> <td></td> </tr> </tbody> </table> | Value | Name | Description | Project | 0 | SIMD8 | 8 LSBs of the execution mask are used | | 1 | SIMD16 | 16 LSBs used in execution mask | | 2 | SIMD32 | 32 bits of execution mask used | |
|-------|--------|---|---------|------|-------------|---------|---|-------|---------------------------------------|--|---|--------|--------------------------------|--|---|--------|--------------------------------|--|
| Value | Name | Description | Project | | | | | | | | | | | | | | | |
| 0 | SIMD8 | 8 LSBs of the execution mask are used | | | | | | | | | | | | | | | | |
| 1 | SIMD16 | 16 LSBs used in execution mask | | | | | | | | | | | | | | | | |
| 2 | SIMD32 | 32 bits of execution mask used | | | | | | | | | | | | | | | | |
| | 29:22 | <p>Reserved Project: All Format: MBZ</p> | | | | | | | | | | | | | | | | |
| | 21:16 | <p>Thread Depth Counter Maximum</p> <p>The maximum value of the thread depth counter.</p> <p>Thread_Depth_Max*Thread_Height_Max*Thread_Width_Max <= 32 for SIMD32, <= 64 for SIMD16 and SIMD8</p> | | | | | | | | | | | | | | | | |
| | 15:14 | <p>Reserved Project: All Format: MBZ</p> | | | | | | | | | | | | | | | | |
| | 13:8 | <p>Thread Height Counter Maximum: The maximum value of the thread height counter</p> | | | | | | | | | | | | | | | | |
| | 7:6 | <p>Reserved Project: All Format: MBZ</p> | | | | | | | | | | | | | | | | |
| | 5:0 | <p>Thread Width Counter Maximum: The maximum value of the thread width counter</p> | | | | | | | | | | | | | | | | |
| 3 | 31:0 | <p>Thread Group ID Starting X</p> <p>This is the initial value of the X component of the thread group</p> | | | | | | | | | | | | | | | | |
| 4 | 31:0 | <p>Thead Group ID X Dimension</p> <p>The X dimension of the thread group (maximum X is dimension -1)</p> | | | | | | | | | | | | | | | | |
| 5 | 31:0 | <p>Thread Group ID Starting Y</p> <p>This is the initial value of the Y component of the thread group</p> | | | | | | | | | | | | | | | | |
| 6 | 31:0 | <p>Thread Group ID Y Dimension</p> <p>The Y dimension of the thread group (maximum Y is dimension -1)</p> | | | | | | | | | | | | | | | | |
| 7 | 31:0 | <p>Thread Group ID Starting Z</p> <p>This is the initial value of the Z component of the thread group</p> | | | | | | | | | | | | | | | | |
| 8 | 31:0 | <p>Thead Group ID Z Dimension</p> <p>The Z dimension of the thread group (maximum Z is dimension -1)</p> | | | | | | | | | | | | | | | | |
| 9 | 31:0 | <p>Right Execution Mask</p> | | | | | | | | | | | | | | | | |
| 10 | 31:0 | <p>Bottom Execution Mask</p> | | | | | | | | | | | | | | | | |

1.8 Media Messages

All message formats are given in terms of dwords (32 bits) using the following conventions which are detailed in GEN4 Subsystem Chapter.

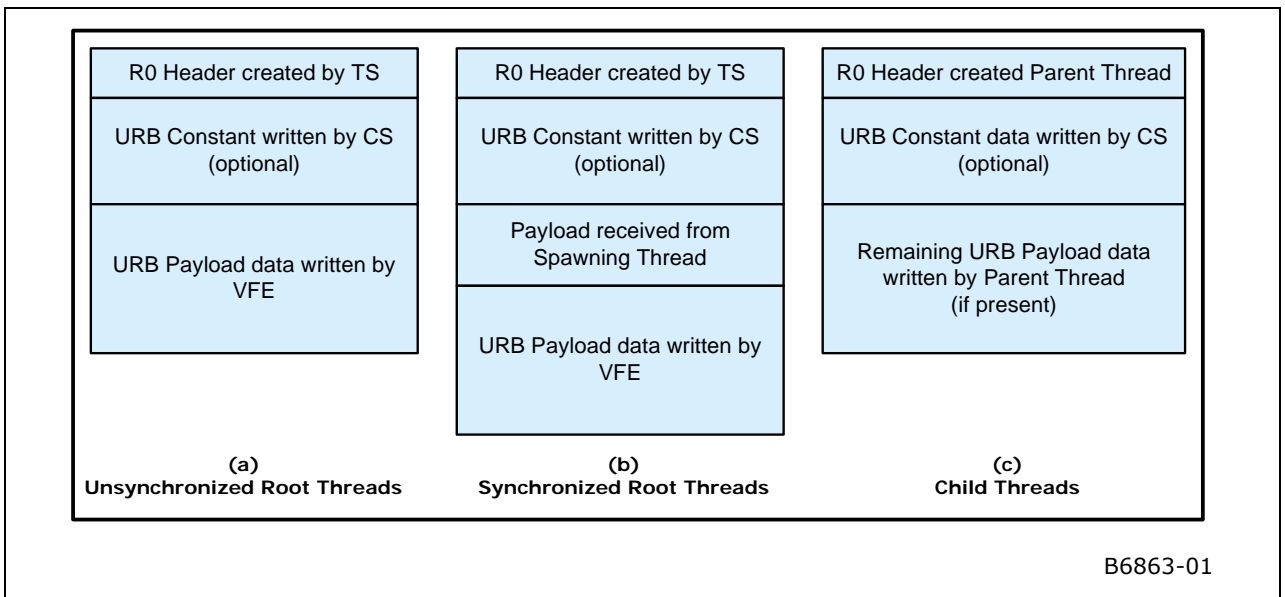
Dispatch Messages: **Rp.d**

SEND Instruction Messages: **Mp.d**

1.8.1 Thread Payload Messages

The root thread's register contents differ from that of child threads, as shown in Figure 1-5. The register contents for a synchronized root thread (also referred to as 'spawned root thread') and an unsynchronized one are also different. Whether the URB Constant data field is present or not is determined by the interface descriptor of a given thread. This applies to both root and child threads. When URB Constant data field is present for a synchronized root thread, URB constant data field is before the data field received from the spawning thread, which is also before the URB payload data.

Figure 1-5. Thread payload message formats for root and child threads



1.8.1.1 Generic Mode Root Thread

The following table shows the R0 register contents for a Generic mode root thread, which is generated by TS. The remaining payloads are application dependent.



R0 header of a generic mode root thread

| DWord | Bit | Description |
|-------------------|-------|---|
| R0.7 | 31 | Reserved |
| | 27:24 | Reserved |
| | 23:0 | Reserved |
| R0.6 | 31:24 | Reserved |
| | 23:0 | Reserved |
| R0.5 | 31:10 | Scratch Space Pointer. Specifies the 1k-byte aligned pointer to the scratch space. This field is only valid when Scratch Space is enabled. Format = GeneralStateOffset[31:10] |
| | 9:8 | Reserved : MBZ |
| | 7:0 | FTID. This ID is assigned by TS and is a unique identifier for the thread in comparison to other concurrent root threads. It is used to free up resources used by the thread upon thread completion. |
| R0.4 | 31:5 | Binding Table Pointer: Specifies the 32-byte aligned pointer to the Binding Table. It is specified as an offset from the Surface State Base Address . Format = SurfaceStateOffset[31:5] |
| | 4:0 | Reserved : MBZ |
| R0.3 | 31:5 | Sampler State Pointer. Specifies the 32-byte aligned pointer to the sampler state table. Format = GeneralStateOffset[31:5] |
| | 4 | Reserved : MBZ |
| | 3:0 | Per Thread Scratch Space. Specifies the amount of scratch space, in 16-byte quantities, allowed to be used by this thread. The value specifies the power that two will be raised to, to determine the amount of scratch space. Format = U4 Range = [0,11] indicating [1k bytes, 2M bytes] in powers of two |
| R0.2 [Pre-DevSNB] | 31:4 | Interface Descriptor Pointer. Specifies the 16-byte aligned pointer to <i>this thread's</i> interface descriptor. Can be used as a base from which to offset child thread's interface descriptor pointers from. Format = GeneralStateOffset[31:4] |
| | 3:0 | Reserved : MBZ |
| | 27:24 | BarrierID. This field indicates which one from the 16 Barriers this kernel is associated. Format: U4 |
| | 23:21 | Reserved : MBZ |
| | 20:16 | Interface Descriptor Offset. This field specifies the offset from the interface descriptor base pointer to the interface descriptor which will be applied to this object. It is specified in units of interface descriptors. Format = U5 |
| | 15:4 | Reserved : MBZ |



| DWord | Bit | Description |
|-------|-------|---|
| | 3:0 | Scoreboard Color (only with MEDIA_OBJECT_EX): This field specifies which dependency color the current thread belongs to. It affects the dependency scoreboard control. Format = U4 |
| R0.1 | 31:28 | [DevILK] Scoreboard Mask 4-7 (only with MEDIA_OBJECT_EX): Each bit indicates the corresponding dependency scoreboard is dependent on. This field is AND'd with the corresponding Scoreboard Mask field in the VFE_STATE_EX. Bit n (for n = 4...7): Scoreboard n is dependent, where bit 28 maps to n = 4. Format = TRUE/FALSE [Pre-DevILK] Reserved : MBZ |
| | 27:26 | [DevILK] Scoreboard Color (only with MEDIA_OBJECT_EX): This field specifies which dependency color the current thread belongs to. It affects the dependency scoreboard control. Format = U2 [Pre-DevILK] Reserved : MBZ |
| | 25 | Reserved. MBZ |
| | 24:16 | [DevILK+] Scoreboard Y This field provides the Y term of the scoreboard value of the current thread. Format = U9 [DevILK] only with MEDIA_OBJECT_EX. [Pre-DevILK] Reserved : MBZ |
| | 15:12 | [DevILK] Scoreboard Mask 0-3 (only with MEDIA_OBJECT_EX): Each bit indicates the corresponding dependency scoreboard is dependent on. This field is AND'd with the corresponding Scoreboard Mask field in the VFE_STATE_EX. Bit n (for n = 0...3): Scoreboard n is dependent, where bit 12 maps to n = 0. Format = TRUE/FALSE [Pre-DevILK] Reserved : MBZ |
| | 11:9 | Reserved. MBZ |
| | 8:0 | [DevILK+] Scoreboard X This field provides the X term of the scoreboard value of the current thread. Format = U9 [DevILK] only with MEDIA_OBJECT_EX. [Pre-DevILK] Reserved : MBZ |
| | 23:16 | [DevILK+] Thread Dependency Identifier (TDID). This field is assigned by TS to be used for hardware scoreboard. Format = U8 [Pre-DevILK] Reserved : MBZ |
| | 15:0 | URB Handle. This is the URB handle where indicating the URB space for use by the root thread and its children. |



1.8.1.2 Root Thread from MEDIA_OBJECT_PRT [DevCTG+]

The root thread payload message for an MEDIA_OBJECT_PRT command has a fixed format independent of the VFE mode (e.g. Generic mode or AVC-IT mode). One example GRF register location is given for the condition that CURBE is disabled.

Root thread payload layout for a MEDIA_OBJECT_PRT command

| GRF Register | Example | Description |
|--------------|---------|--|
| R0 | R0 | R0 header |
| R1 – R(m) | n/a | Constants from CURBE when CURBE is enabled m is a non-negative value |
| R(m+1) | R1 | In-line Data block. |

The R0 header field is as the following, which is the same as in other modes except the Thread Restart Enable bit (bit 0 of R0.2).

The inline data block field is the same as in the MEDIA_OBJECT_EX command with zero-filled partial GRF.

1.8.1.3 IS-Mode Root Thread [Pre-DevSNB]

The following table shows the root thread payload messages when VFE is in IS mode and URB push constant is not enabled.

| DWord | Bit | Description |
|-------|-------|---|
| R0.7 | 31 | Reserved |
| | 27:24 | Reserved |
| | 23:0 | Reserved : MBZ. |
| R0.6 | 31:24 | Reserved |
| | 23:0 | Reserved |
| R0.5 | 31:10 | NOT USED (was Scratch Space Pointer). |
| | 9:8 | Reserved : MBZ |
| | 7:0 | FFTID . This ID is assigned by TS and is a unique identifier for the thread in comparison to other concurrent root threads. It is used to free up resources used by the thread upon thread completion. Note: Nothing to free up in this case. |
| R0.4 | 31:5 | Binding Table Pointer : Specifies the 32-byte aligned pointer to the Binding Table. It is specified as an offset from the Surface State Base Address . Format = SurfaceStateOffset[31:5] |
| | 4:0 | Reserved : MBZ |
| R0.3 | 31:5 | NOT USED (was Sampler State Pointer). |



| DWord | Bit | Description |
|-------|-------|---|
| | 4 | Reserved : MBZ |
| | 3:0 | NOT USED (was Per Thread Scratch Space) |
| R0.2 | 31:5 | Interface Descriptor Pointer. Specifies the 32-byte aligned pointer to <i>this thread's</i> interface descriptor. Can be used as a base from which to offset child thread's interface descriptor pointers from. Format = GeneralStateOffset[31:5] |
| | 4:0 | Reserved : MBZ |
| R0.1 | 31:0 | Reserved : MBZ |
| R0.0 | 31:16 | Reserved : MBZ |
| | 15:0 | URB Handle. This is the URB handle where indicating the URB space for use by the root thread and its children. This may be used if child threads and/or synchronized root threads are present in IS mode. |
| R1.7 | 31:16 | Motion Vectors – Field 1, Backward, Vertical Component. Each vector component is a 16-bit two's-complement value. The vector is relative to the current macroblock location. According to ISO/IEC 13818-2 Table 7-8, the valid range of each vector component is [-2048, +2047.5], implying a format of s11.1. However, it should be noted that motion vector values are sign extended to 16 bits. |
| | 15:0 | Motion Vectors – Field 1, Backward, Horizontal Component |
| R1.6 | 31:16 | Motion Vectors – Field 1, Forward, Vertical Component |
| | 15:0 | Motion Vectors – Field 1, Forward, Horizontal Component |
| R1.5 | 31:16 | Motion Vectors – Field 0, Backward, Vertical Component |
| | 15:0 | Motion Vectors – Field 0, Backward, Horizontal Component |
| R1.4 | 31:16 | Motion Vectors – Field 0, Forward, Vertical Component |
| | 15:0 | Motion Vectors – Field 0, Forward, Horizontal Component |
| R1.3 | 31:24 | Subblock Coding for Block Cr5 |
| | 23:16 | Subblock Coding for Block Cb4 |
| | 15:8 | Subblock Coding for Block Y3 |
| | 7:0 | Subblock Coding for Block Y2 |
| R1.2 | 31:24 | Subblock Coding for Block Y1 |
| | 23:16 | Subblock Coding for Block Y0. This field specifies the subblock partition and subblock coding pattern for the block. The definition of the 8 bits of this field is listed below. Detailed coding can be found in Error! Reference source not found.. Bits [7:6]: reserved Bits [5:2]: Subblock present Bits [1:0]: Subblock partitioning |
| | 15:12 | Reserved. |



| DWord | Bit | Description | | | | | | | | | | | | | | | | | | | | |
|-------|--|---|---|--|-------------|---------------------------------|----------|------|-------|-------|------|-------|------|------|------------|------------|---|---|----|---|---|---|
| | 11:6 | <p>Coded Block Pattern. This field specifies whether blocks are present or not. Format = 6-bit mask. Bit 11: Y0 Bit 10: Y1 Bit 9: Y2 Bit 8: Y3 Bit 7: Cb5 Bit 6: Cr5</p> | | | | | | | | | | | | | | | | | | | | |
| | 5:0 | Reserved. | | | | | | | | | | | | | | | | | | | | |
| R1.1 | 31:24 | Reserved. (Skip Macroblocks) | | | | | | | | | | | | | | | | | | | | |
| | 23:0 | Reserved. (Offset into error data) | | | | | | | | | | | | | | | | | | | | |
| R1.0 | 31:28 | <p>Motion Vertical Field Select. A bit-wise representation of a long [2][2] array as defined in §6.3.17.2 of the <i>ISO/IEC 13818-2</i> (see also §7.6.4).</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>MVector [r]</th> <th>MVector [s]</th> <th>MotionVerticalFieldSelect Index</th> </tr> </thead> <tbody> <tr> <td>28</td> <td>0</td> <td>0</td> <td>0</td> </tr> <tr> <td>29</td> <td>0</td> <td>1</td> <td>1</td> </tr> <tr> <td>30</td> <td>1</td> <td>0</td> <td>2</td> </tr> <tr> <td>31</td> <td>1</td> <td>1</td> <td>3</td> </tr> </tbody> </table> <p>Format = MC_MotionVerticalFieldSelect. 0 = The prediction is taken from the <u>top</u> reference field. 1 = The prediction is taken from the <u>bottom</u> reference field.</p> | Bit | MVector [r] | MVector [s] | MotionVerticalFieldSelect Index | 28 | 0 | 0 | 0 | 29 | 0 | 1 | 1 | 30 | 1 | 0 | 2 | 31 | 1 | 1 | 3 |
| | Bit | MVector [r] | MVector [s] | MotionVerticalFieldSelect Index | | | | | | | | | | | | | | | | | | |
| | 28 | 0 | 0 | 0 | | | | | | | | | | | | | | | | | | |
| | 29 | 0 | 1 | 1 | | | | | | | | | | | | | | | | | | |
| 30 | 1 | 0 | 2 | | | | | | | | | | | | | | | | | | | |
| 31 | 1 | 1 | 3 | | | | | | | | | | | | | | | | | | | |
| 27 | <p>Second Field. This bit indicates that this is the second field in the current frame. The prediction for this macroblock, if it belongs to a field P-picture, should use this bit to determine which frame contains the reference field as described in §7.6.2.1 of the <i>ISO/IEC 13818-2</i>.</p> <p>When the picture type is not P or the prediction type is not field, this bit is set to 0.</p> <p>Format = MC_SecondPField 0 = This is not the second field. 1 = This is the second field.</p> | | | | | | | | | | | | | | | | | | | | | |
| 26 | Reserved. (HWMC mode) | | | | | | | | | | | | | | | | | | | | | |
| 25:24 | <p>Motion Type. When combined with the destination picture type (field or frame) this Motion Type field indicates the type of motion to be applied to the macroblock. See <i>ISO/IEC 13818-2</i> §6.3.17.1, Tables 6-17, 6-18. In particular, the device supports dual-prime motion prediction (11) in both frame and field picture type.</p> <p>Format = MC_MotionType</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Destination = Frame Picture_Structure = 11</th> <th>Destination = Field Picture_Structure != 11</th> </tr> </thead> <tbody> <tr> <td>'00'</td> <td>Reserved</td> <td>Reserved</td> </tr> <tr> <td>'01'</td> <td>Field</td> <td>Field</td> </tr> <tr> <td>'10'</td> <td>Frame</td> <td>16x8</td> </tr> <tr> <td>'11'</td> <td>Dual-Prime</td> <td>Dual-Prime</td> </tr> </tbody> </table> | Value | Destination = Frame Picture_Structure = 11 | Destination = Field Picture_Structure != 11 | '00' | Reserved | Reserved | '01' | Field | Field | '10' | Frame | 16x8 | '11' | Dual-Prime | Dual-Prime | | | | | | |
| Value | Destination = Frame Picture_Structure = 11 | Destination = Field Picture_Structure != 11 | | | | | | | | | | | | | | | | | | | | |
| '00' | Reserved | Reserved | | | | | | | | | | | | | | | | | | | | |
| '01' | Field | Field | | | | | | | | | | | | | | | | | | | | |
| '10' | Frame | 16x8 | | | | | | | | | | | | | | | | | | | | |
| '11' | Dual-Prime | Dual-Prime | | | | | | | | | | | | | | | | | | | | |



| DWord | Bit | Description |
|-------|-------|---|
| | 23:22 | Reserved. (Scan method) |
| | 21 | DCT Type. This field specifies the DCT type of the current macroblock. The kernel should ignore this field when processing Cb/Cr data. See <i>ISO/IEC 13818-2</i> §6.3.17.1. This field is zero if Coded Block Pattern is also zero (no coded blocks present). 0 = MC_FRAME_DCT (Macroblock is frame DCT coded). 1 = MC_FIELD_DCT (Macroblock is field DCT coded). |
| | 20 | Reserved. (H261 Loop Filter) |
| | 19 | Reserved. (H263) |
| | 18 | Macroblock Motion Backward. This field specifies if the backward motion vector is active. See <i>ISO/IEC 13818-2</i> Tables B-2 through B-4. 0 = No backward motion vector. 1 = Use backward motion vector(s). |
| | 17 | Macroblock Motion Forward. This field specifies if the forward motion vector is active. See <i>ISO/IEC 13818-2</i> Tables B-2 through B-4. 0 = No forward motion vector. 1 = Use forward motion vector(s). |
| | 16 | Macroblock Intra Type. This field specifies if the current macroblock is intra-coded. When set, Coded Block Pattern is ignored and no prediction is performed (i.e., no motion vectors are used). See <i>ISO/IEC 13818-2</i> Tables B-2 through B-4. 0 = Non-intra macroblock. 1 = Intra macroblock. |
| | 15:0 | Reserved. |
| R2.7 | 31:0 | Reserved. |
| R2.6 | 31:0 | Reserved. |
| R2.5 | 31:0 | Reserved. |
| R2.4 | 31:0 | Reserved. |
| R2.3 | 31:0 | Reserved. |
| R2.2 | 31:0 | Reserved. |
| R2.1 | 31:27 | Reserved. |
| | 26:20 | Vertical Origin. Set the vertical origin of the next macroblock in the destination picture in units of macroblocks. (Valid range is 0 to 120). Format = U7 in macroblock units. Range = [0, 120] |
| | 19:11 | Reserved: MBZ |
| | 10:4 | Horizontal Origin. Set the horizontal origin of the next macroblock in the destination picture in units of macroblocks. Format = U7 in macroblock units. Range = [0, 127] |
| | 3:0 | Reserved. |
| R2.0 | 31:30 | Reserved. |



| DWord | Bit | Description |
|--|-------|---|
| | 29 | May need this for WMV. (Interpolation Rounder Control) |
| | 28 | May need this for WMV. (Bidirectional Averaging Control) |
| | 27:20 | Reserved. |
| | 19:18 | Picture Coding Type??? . This field identifies whether the picture is an intra-coded picture (I), predictive-coded picture (P) or bi-directionally predictive-coded picture (B). See <i>ISO/IEC 13818-2</i> §6.3.9 for details. Format = MPEG_PICTURE_CODING_TYPE 00 = Reserved 01 = MPEG_I_PICTURE 10 = MPEG_P_PICTURE 11 = MPEG_B_PICTURE |
| | 17:16 | Picture Structure??? . This field specifies whether the picture is encoded in the form of a frame picture or one field (top or bottom) picture. See <i>ISO/IEC 13818-2</i> §6.3.10 for details. Format = MPEG_PICTURE_STRUCTURE 00 = Reserved 01 = MPEG_TOP_FIELD 10 = MPEG_BOTTOM_FIELD 11 = MPEG_FRAME |
| | 15 | Reserved. (8-bit Intra) |
| | 14:13 | Reserved. (Intra DC Precision) |
| | 12:0 | Reserved. |
| None (0 blocks coded) or R3- R[2+4x] where x = number of coded blocks | | DCT Coefficients. These are the DCT values of the coefficients for the macroblock. Only coded blocks have coefficients present in the array. Beginning in R3, the order of the coefficients for the coded blocks is Y0, Y1, Y2, Y3, Cb4, and Cr5. For each coded block, the 8x8 DCT coefficients, with 1 word each coefficient, are organized in row-major order, occupying four GRF registers. This is shown in Table 1-1, where the index-pair for a DCT coefficient is (Column_Index, Row_Index). |



1.8.1.4 VLD-Mode Root Thread [Pre-DevSNB]

The following table shows the root thread payload messages when VFE is in VLD mode and URB push constant is not enabled. When URB push constant is enabled, it will start at R1. Subsequently, macroblock data starting with motion vectors will be put in GRF registers after the URB push constants.

| DWord | Bit | Description |
|-------|-------|---|
| R0.7 | 31 | Reserved |
| | 27:24 | Reserved |
| | 23:0 | Reserved : MBZ. |
| R0.6 | 31:24 | Reserved |
| | 23:0 | Reserved |
| R0.5 | 31:10 | NOT USED (was Scratch Space Pointer). |
| | 9:8 | Reserved : MBZ |
| | 7:0 | FFTID. This ID is assigned by TS and is a unique identifier for the thread in comparison to other concurrent root threads. It is used to free up resources used by the thread upon thread completion. Note: Nothing to free up in this case. |
| R0.4 | 31:5 | Binding Table Pointer: Specifies the 32-byte aligned pointer to the Binding Table. It is specified as an offset from the Surface State Base Address . Format = SurfaceStateOffset[31:5] |
| | 4:0 | Reserved : MBZ |
| R0.3 | 31:5 | NOT USED (was Sampler State Pointer). |
| | 4 | Reserved : MBZ |
| | 3:0 | NOT USED (was Per Thread Scratch Space) |
| R0.2 | 31:5 | Interface Descriptor Pointer. Specifies the 32-byte aligned pointer to <i>this thread's</i> interface descriptor. Can be used as a base from which to offset child thread's interface descriptor pointers from. Format = GeneralStateOffset[31:5] |
| | 4:0 | Reserved : MBZ |
| R0.1 | 31:0 | Reserved : MBZ |
| R0.0 | 31:16 | Reserved : MBZ |
| | 15:0 | NOT USED (was URB Handle) |
| R1.7 | 31:16 | Motion Vectors – Field 1, Backward, Vertical Component. Each vector component is a 16-bit two's-complement value. The vector is relative to the current macroblock location. According to ISO/IEC 13818-2 Table 7-8, the valid range of each vector component is [-2048, +2047.5], implying a format of s11.1. However, it should be noted that motion vector values are sign extended to 16 bits. |
| | 15:0 | Motion Vectors – Field 1, Backward, Horizontal Component |
| R1.6 | 31:16 | Motion Vectors – Field 1, Forward, Vertical Component |



| DWord | Bit | Description |
|-------|---|---|
| | 15:0 | Motion Vectors – Field 1, Forward, Horizontal Component |
| R1.5 | 31:16 | Motion Vectors – Field 0, Backward, Vertical Component |
| | 15:0 | Motion Vectors – Field 0, Backward, Horizontal Component |
| R1.4 | 31:16 | Motion Vectors – Field 0, Forward, Vertical Component |
| | 15:0 | Motion Vectors – Field 0, Forward, Horizontal Component |
| R1.3 | 31:27 | Reserved. |
| | 26:20 | Vertical Origin. Set the vertical origin of the next macroblock in the destination picture in units of macroblocks. (Valid range is 0 to 120). Format = U7 in macroblock units. Range = [0, 120] |
| | 19:11 | Reserved: MBZ |
| | 10:4 | Horizontal Origin. Set the horizontal origin of the next macroblock in the destination picture in units of macroblocks. Format = U7 in macroblock units. Range = [0, 127] |
| | 3:0 | Reserved. |
| R1.2 | 31:30 | Reserved. |
| | 29 | Reserved. (Interpolation Rounder Control) |
| | 28 | Reserved. (Bidirectional Averaging Control) |
| | 27:20 | Reserved. |
| | 19:18 | Picture Coding Type. This field identifies whether the picture is an intra-coded picture (I), predictive-coded picture (P) or bi-directionally predictive-coded picture (B). See <i>ISO/IEC 13818-2</i> §6.3.9 for details. Format = MPEG_PICTURE_CODING_TYPE 00 = Reserved 01 = MPEG_I_PICTURE 10 = MPEG_P_PICTURE 11 = MPEG_B_PICTURE |
| | 17:16 | Picture Structure. This field specifies whether the picture is encoded in the form of a frame picture or one field (top or bottom) picture. See <i>ISO/IEC 13818-2</i> §6.3.10 for details. Format = MPEG_PICTURE_STRUCTURE 00 = Reserved 01 = MPEG_TOP_FIELD 10 = MPEG_BOTTOM_FIELD 11 = MPEG_FRAME |
| | 15 | Reserved. (8-bit Intra) |
| | 14:13 | Intra DC Precision. See <i>ISO/IEC 13818-2</i> §6.3.10 for details. |
| 12 | Disable Mismatch. This bit is used to disable the mismatch control performed after the inverse quantization operation, as described in <i>ISO/IEC 13818-2</i> §7.4.4 | |



| DWord | Bit | Description | | | | | | | | | | | | | | | | | | | | |
|-------|---|---|-------------|---------------------------------|-------------|---------------------------------|----|---|---|---|----|---|---|---|----|---|---|---|----|---|---|---|
| | 11:6 | <p>Coded Block Pattern. This field specifies whether blocks are present or not. Format = 6-bit mask. Bit 11: Y0 Bit 10: Y1 Bit 9: Y2 Bit 8: Y3 Bit 7: Cb5 Bit 6: Cr5</p> | | | | | | | | | | | | | | | | | | | | |
| | 5 | <p>Quantizer Scale Type: This field specifies the quantizer scaling type. Format = MPEG_Q_SCALE_TYPE 0: MPEG_QSCALE_LINEAR 1: MPEG_QSCALE_NONLINEAR</p> | | | | | | | | | | | | | | | | | | | | |
| | 4:0 | <p>Quantization Scale Code. Combined with the quantization scale type, this value selects the quantizer scale table according to ISO/IEC 13818-2 Table 7-6</p> | | | | | | | | | | | | | | | | | | | | |
| R1.1 | 31:24 | Reserved. (Skip Macroblocks) | | | | | | | | | | | | | | | | | | | | |
| | 23:0 | Reserved. (Offset into error data) | | | | | | | | | | | | | | | | | | | | |
| R1.0 | 31:28 | <p>Motion Vertical Field Select. A bit-wise representation of a long [2][2] array as defined in §6.3.17.2 of the <i>ISO/IEC 13818-2</i> (see also §7.6.4).</p> <table border="1" data-bbox="631 989 1344 1146"> <thead> <tr> <th>Bit</th> <th>MVector [r]</th> <th>MVector [s]</th> <th>MotionVerticalFieldSelect Index</th> </tr> </thead> <tbody> <tr> <td>28</td> <td>0</td> <td>0</td> <td>0</td> </tr> <tr> <td>29</td> <td>0</td> <td>1</td> <td>1</td> </tr> <tr> <td>30</td> <td>1</td> <td>0</td> <td>2</td> </tr> <tr> <td>31</td> <td>1</td> <td>1</td> <td>3</td> </tr> </tbody> </table> <p>Format = MC_MotionVerticalFieldSelect. 0 = The prediction is taken from the <u>top</u> reference field. 1 = The prediction is taken from the <u>bottom</u> reference field.</p> | Bit | MVector [r] | MVector [s] | MotionVerticalFieldSelect Index | 28 | 0 | 0 | 0 | 29 | 0 | 1 | 1 | 30 | 1 | 0 | 2 | 31 | 1 | 1 | 3 |
| | Bit | MVector [r] | MVector [s] | MotionVerticalFieldSelect Index | | | | | | | | | | | | | | | | | | |
| | 28 | 0 | 0 | 0 | | | | | | | | | | | | | | | | | | |
| 29 | 0 | 1 | 1 | | | | | | | | | | | | | | | | | | | |
| 30 | 1 | 0 | 2 | | | | | | | | | | | | | | | | | | | |
| 31 | 1 | 1 | 3 | | | | | | | | | | | | | | | | | | | |
| 27 | <p>Second Field. This bit indicates that this is the second field in the current frame. The prediction for this macroblock, if it belongs to a field P-picture, should use this bit to determine which frame contains the reference field as described in §7.6.2.1 of the <i>ISO/IEC 13818-2</i>.</p> <p>When the picture type is not P or the prediction type is not field, this bit is set to 0.</p> <p>Format = MC_SecondPField 0 = This is not the second field. 1 = This is the second field.</p> | | | | | | | | | | | | | | | | | | | | | |
| 26 | Reserved. (HWMC mode) | | | | | | | | | | | | | | | | | | | | | |



| DWord | Bit | Description | | | | | | | | | | | | | | | |
|-------|---|---|-------|---|--|------|----------|----------|------|-------|-------|------|-------|------|------|------------|------------|
| | 25:24 | <p>Motion Type. When combined with the destination picture type (field or frame) this Motion Type field indicates the type of motion to be applied to the macroblock. See <i>ISO/IEC 13818-2</i> §6.3.17.1, Tables 6-17, 6-18. In particular, the device supports dual-prime motion prediction (11) in both frame and field picture type.</p> <p>Format = MC_MotionType</p> <table border="1" data-bbox="623 506 1357 695"> <thead> <tr> <th>Value</th> <th>Destination = Frame Picture_Structure = 11</th> <th>Destination = Field Picture_Structure != 11</th> </tr> </thead> <tbody> <tr> <td>'00'</td> <td>Reserved</td> <td>Reserved</td> </tr> <tr> <td>'01'</td> <td>Field</td> <td>Field</td> </tr> <tr> <td>'10'</td> <td>Frame</td> <td>16x8</td> </tr> <tr> <td>'11'</td> <td>Dual-Prime</td> <td>Dual-Prime</td> </tr> </tbody> </table> | Value | Destination = Frame Picture_Structure = 11 | Destination = Field Picture_Structure != 11 | '00' | Reserved | Reserved | '01' | Field | Field | '10' | Frame | 16x8 | '11' | Dual-Prime | Dual-Prime |
| Value | Destination = Frame Picture_Structure = 11 | Destination = Field Picture_Structure != 11 | | | | | | | | | | | | | | | |
| '00' | Reserved | Reserved | | | | | | | | | | | | | | | |
| '01' | Field | Field | | | | | | | | | | | | | | | |
| '10' | Frame | 16x8 | | | | | | | | | | | | | | | |
| '11' | Dual-Prime | Dual-Prime | | | | | | | | | | | | | | | |
| | 23:22 | Reserved. (Scan method) | | | | | | | | | | | | | | | |
| | 21 | <p>DCT Type. This field specifies the DCT type of the current macroblock. The kernel should ignore this field when processing Cb/Cr data. See <i>ISO/IEC 13818-2</i> §6.3.17.1. This field is zero if Coded Block Pattern is also zero (no coded blocks present).</p> <p>0 = MC_FRAME_DCT (Macroblock is frame DCT coded). 1 = MC_FIELD_DCT (Macroblock is field DCT coded).</p> | | | | | | | | | | | | | | | |
| | 20 | Reserved. (H261 Loop Filter) | | | | | | | | | | | | | | | |
| | 19 | Reserved. (H263) | | | | | | | | | | | | | | | |
| | 18 | <p>Macroblock Motion Backward. This field specifies if the backward motion vector is active. See <i>ISO/IEC 13818-2</i> Tables B-2 through B-4.</p> <p>0 = No backward motion vector. 1 = Use backward motion vector(s).</p> | | | | | | | | | | | | | | | |
| | 17 | <p>Macroblock Motion Forward. This field specifies if the forward motion vector is active. See <i>ISO/IEC 13818-2</i> Tables B-2 through B-4.</p> <p>0 = No forward motion vector. 1 = Use forward motion vector(s).</p> | | | | | | | | | | | | | | | |
| | 16 | <p>Macroblock Intra Type. This field specifies if the current macroblock is intra-coded. When set, Coded Block Pattern is ignored and no prediction is performed (i.e., no motion vectors are used). See <i>ISO/IEC 13818-2</i> Tables B-2 through B-4.</p> <p>0 = Non-intra macroblock. 1 = Intra macroblock.</p> | | | | | | | | | | | | | | | |
| | 15:0 | Reserved. | | | | | | | | | | | | | | | |



| DWord | Bit | Description |
|---|-----|---|
| None (0 block coded) or R2- R[1+4x] where x = number of coded blocks | | DCT Coefficients. These are the DCT values of the coefficients for the macroblock. Only coded blocks have coefficients present in the array. Beginning in R2, the order of the coefficients for the coded blocks is Y0, Y1, Y2, Y3, Cb4, and Cr5. For each coded block, the 8x8 DCT coefficients, with 1 word each coefficient, are organized in row-major order, occupying four GRF registers. This is shown in Table 1-1, where the index-pair for a DCT coefficient is (Column_Index, Row_Index). |

Table 1-1. Format of a block of DCT coefficients in GRF registers

| Reg. / Words | W15 | W14 | W13 | W12 | W11 | W10 | W9 | W8 | W7 | W6 | W5 | W4 | W3 | W2 | W1 | W0 |
|--------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-----|-----|-----|-----|-------|-------|
| R[n] | (7,1) | (6,1) | (5,1) | (4,1) | (3,1) | (2,1) | (1,1) | (0,1) | (7,0) | (6,0) | ... | ... | ... | ... | (1,0) | (0,0) |
| R[n+1] | (7,3) | ... | ... | ... | ... | ... | ... | (0,3) | (7,2) | (6,2) | ... | ... | ... | ... | (1,2) | (0,2) |
| R[n+2] | (7,5) | ... | ... | ... | ... | ... | ... | (0,5) | (7,4) | (6,4) | ... | ... | ... | ... | (1,4) | (0,4) |
| R[n+3] | (7,7) | (7,7) | (5,7) | (4,7) | (3,7) | (2,7) | (1,7) | (0,7) | (7,6) | (6,6) | ... | ... | ... | ... | (1,6) | (0,6) |

* W# (# from 0 to 15) represents WORD location # within an 8-DW register.



1.8.1.5 AVC-IT Mode Root Thread [DevCTG, DevILK]

The following table shows the root thread payload messages when VFE is in AVC-IT mode. One example GRF register location is given for the condition that: CURBE is disabled, Weight-Offset Offset (p) = 4 and Residual Offset (q) = 6.

AVC-IT Mode root thread payload layout

| GRF Register | Example | Description |
|----------------------|----------|---|
| R0 | R0 | R0 header |
| R1 – R(m) | n/a | Constants from CURBE when CURBE is enabled m is a non-negative value. If m is 0, this field doesn't exist. |
| R(m+1) | R1 | In-line Data block. |
| R(m+2) – R(m+2+n-1) | R2 – R5 | Indirect Data Motion Vector blocks If maximum MvSize is 32, the number of GRF space occupied by motion vector, n, is 4. Otherwise (if maximum MvSize is 16), n is 2. |
| R(m+p+2) – R(m+p+3) | R6 – R7 | Indirect Data Weight-Offset blocks Weight-Offset Offset (p) must be greater than or equal to n. When p is greater than n, the gap, which is R(m+2+n) to R(m+p+1), contains undefined values. |
| R(m+q+2) – R(m+q+13) | R8 – R31 | Indirect Data Residual blocks Residual Offset (q) must be greater than or equal to p+2. When q is greater than p+2, the gap, which is R(m+p+4) to R(m+q+1), contains undefined values. |

The R0 header field is the same as in other modes.

AVC-IT Mode R0 Header

| DWord | Bit | Description |
|-------|-------|---|
| R0.0 | 31:24 | Reserved : MBZ |
| | 23:16 | [DevILK]: Thread Dependency Identifier (TDID). This field is assigned by TS to be used for hardware scoreboard. Format = U8 [DevCTG] Reserved : MBZ |
| | 15:0 | URB Handle. This is the URB handle where indicating the URB space for use by the root thread and its children. This may be used if child threads and/or synchronized root threads are present in IS mode. |
| R0.1 | 31:28 | [DevILK] Scoreboard Mask 4-7 (only with MEDIA_OBJECT_EX): Each bit indicates the corresponding dependency scoreboard is dependent on. This field is AND'd with the corresponding Scoreboard Mask field in the VFE_STATE_EX. Bit n (for n = 4...7): Scoreboard n is dependent, where bit 28 maps to n = 4. Format = TRUE/FALSE [Pre-DevILK] Reserved : MBZ |



| DWord | Bit | Description |
|-------|-------|--|
| | 27:26 | [DevILK] Scoreboard Color (only with MEDIA_OBJECT_EX): This field specifies which dependency color the current thread belongs to. It affects the dependency scoreboard control. Format = U2 [Pre-DevILK] Reserved : MBZ |
| | 25 | Reserved. MBZ |
| | 24:16 | [DevILK] Scoreboard Y (only with MEDIA_OBJECT_EX) This field provides the Y term of the scoreboard value of the current thread. Format = U9 [Pre-DevILK] Reserved : MBZ |
| | 15:12 | [DevILK] Scoreboard Mask 0-3 (only with MEDIA_OBJECT_EX): Each bit indicates the corresponding dependency scoreboard is dependent on. This field is AND'd with the corresponding Scoreboard Mask field in the VFE_STATE_EX. Bit n (for n = 0...3): Scoreboard n is dependent, where bit 12 maps to n = 0. Format = TRUE/FALSE [Pre-DevILK] Reserved : MBZ |
| | 11:9 | Reserved. MBZ |
| | 8:0 | [DevILK] Scoreboard X (only with MEDIA_OBJECT_EX) This field provides the X term of the scoreboard value of the current thread. Format = U9 [Pre-DevILK] Reserved : MBZ |
| R0.2 | 31:5 | Interface Descriptor Pointer. Specifies the 32-byte aligned pointer to <i>this thread's</i> interface descriptor. Can be used as a base from which to offset child thread's interface descriptor pointers from. Format = GeneralStateOffset[31:5] |
| | 4:0 | Reserved : MBZ |
| R0.3 | 31:5 | NOT USED (was Sampler State Pointer). |
| | 4 | Reserved : MBZ |
| | 3:0 | NOT USED (was Per Thread Scratch Space) |
| R0.4 | 31:5 | Binding Table Pointer: Specifies the 32-byte aligned pointer to the Binding Table. It is specified as an offset from the Surface State Base Address . Format = SurfaceStateOffset[31:5] |
| | 4:0 | Reserved : MBZ |
| R0.5 | 31:10 | NOT USED (was Scratch Space Pointer). |
| | 9:8 | Reserved : MBZ |
| | 7:0 | FTTID. This ID is assigned by TS and is a unique identifier for the thread in comparison to other concurrent root threads. It is used to free up resources used by the thread upon thread completion. Note: Nothing to free up in this case. |
| R0.6 | 31:24 | Reserved |
| | 23:0 | Reserved |



| DWord | Bit | Description |
|-------|-------|------------------------|
| R0.7 | 31 | Reserved |
| | 27:24 | Reserved |
| | 23:0 | Reserved : MBZ. |

The in-line data block differs from the in-line data in the MEDIA_OBJECT_EX command by one field – QpPrimeY is replaced by the DerivedMbInfo field.

AVC-IT In-line data block

| DWord | Bit | Description |
|-------|-------|---|
| R+0.0 | 31:28 | Reserved. MBZ |
| | 26:25 | MbAffFieldFlag This field indicates that the current macroblock is a field macroblock within a MbAff frame picture. It is provided as Flag = MbaffFrame & FieldMbFlag . 00 = if (Flag == 0) 11 = if (Flag == 1) Other encodings are reserved |
| | 24 | FieldMbPolarityFlag This field indicates the field polarity of the current macroblock. Within a MbAff frame picture, this field may be different per macroblock and is set to 1 only for the second macroblock in a MbAff pair if FieldMbFlag is set. Otherwise, it is set to 0. Within a field picture, this field is set to 1 if the current picture is the bottom field picture. Otherwise, it is set to 0. It is a constant for the whole field picture. This field is reserved and MBZ for a progressive frame picture. 0 = Current macroblock is a field macroblock from the top field 1 = Current macroblock is a field macroblock from the bottom field <i>Programming Note: Here bits [26:24] (MbAffFieldFlag and FiedMbPolarityFlag) match with bits [10:8] of the Media Block Read message descriptor, simplifying the programming for message generation, as when MbAffFieldFlag is "1", kernels need to override the original "frame" surface state set for MBAFF frame picture.</i> |
| | 23 | Reserved: MBZ |
| | 22:17 | Reserved: MBZ (Forced to zero by VFE to simplify use of this word for data port message) |
| | 16 | Reserved. MBZ |
| | 15 | Transform8x8Flag |
| | 14 | FieldMbFlag (Field Macroblock Flag). This field specifies whether current macroblock is field macroblock. 0 = Frame macroblock. 1 = Field macroblock. |
| | 13 | IntraMacroblock (Intra Macroblock Indicator) |
| | 12:8 | MbType (Macroblock Type). This field, along with "IntraMacroblock" specifies the macroblock types. |



| DWord | Bit | Description |
|-------|-------|---|
| | 7:6 | WeightedBiPredFlag (Weighted Bidirectional Prediction Flag) (from Picture State). Valid only for macroblock in inter mode. Otherwise (intra macroblock), this field is reserved. |
| | 5 | WeightedPredFlag. (from Picture State). Valid only for macroblock in inter mode. Otherwise (intra macroblock), this field is reserved. |
| | 4 | Reserved. MBZ |
| | 3:2 | ChromaFormatIdc (Chroma Format Indicator). This field specifies the chroma format for the decoding process as defined below. This field is constant within a picture. 00 = Luma only (monochrome) 01 = YUV420 sampling 10 is reserved (for YUV422 sampling) 11 is reserved (for YUV444 sampling) |
| | 1 | MbaffFrame. (from Picture State). |
| | 0 | FieldPicFlag. (from Picture State) |
| R+0.1 | 31:16 | CbpY (Coded Block Pattern Y) <i>Not expected to be used by Kernel.</i> |
| | 15:8 | VertOrigin (Vertical Origin). This field specifies the vertical origin of current macroblock in the destination picture in units of macroblocks. For field macroblock pair in MBAFF frame, the vertical origins for both macroblocks should be set as if they were located in corresponding field pictures. For example, for field macroblock pair originated at (16, 64) pixel location in an MBAFF frame picture, the Vertical Origin for both macroblocks should be set as 2 (macroblocks). Format = U8 in unit of macroblock. |
| | 7:0 | HorzOrigin (Horizontal Origin). This field specifies the horizontal origin of current macroblock in the destination picture in units of macroblocks. Format = U8 in unit of macroblock. |
| R+0.2 | 31:14 | Reserved. MBZ (Forced to Zero by VFE) |
| | 13:8 | Coded Block Pattern (VFE Derived). This field specifies whether blocks (8x8) are present or not. Each bit corresponds to one block. "0" indicates error block isn't present, "1" indicates error block is present. Bit 13: Y0 Bit 12: Y1 Bit 11: Y2 Bit 10: Y3 Bit 9: Cb Bit 8: Cr |
| | 7:4 | CbpCr <i>Not expected to be used by Kernel.</i> |
| | 3:0 | CbpCb <i>Not expected to be used by Kernel.</i> |



| DWord | Bit | Description |
|---------------------|--------------|---|
| R+0.3 – R+0.5 | 31:0 each | InterIntraSpecificFields See Error! Reference source not found. for intra macroblocks See Error! Reference source not found. for inter macroblocks |
| R+0.6 | 31:0 | Reserved. MBZ |
| R+0.7 | 31:0 | Reserved. MBZ |

(R+0.3 to R+0.5) for an Intra Macroblock

| Dword | Bit | Description | | | | | | | | | | | | | | | |
|-------|---|--|------|---------------------------------|-----|-----------------------------|---|----------|---|--|---|--------------------------------|---|--------------------------------|---|--------------------------------|---|
| R+0.3 | 31:8 | Reserved | | | | | | | | | | | | | | | |
| | 7:0 | MbIntraStruct (Macroblock Intra Structure) <table border="1" style="margin-left: 20px;"> <thead> <tr> <th>Bits</th> <th>MotionVerticalFieldSelect Index</th> </tr> </thead> <tbody> <tr> <td>7:6</td> <td>ChromalIntraPredMode</td> </tr> <tr> <td>5</td> <td>Reserved</td> </tr> <tr> <td>4</td> <td>IntraPredAvailFlagE – E (Left First Half)</td> </tr> <tr> <td>3</td> <td>IntraPredAvailFlagD – D</td> </tr> <tr> <td>2</td> <td>IntraPredAvailFlagC – C</td> </tr> <tr> <td>1</td> <td>IntraPredAvailFlagB – B</td> </tr> <tr> <td>0</td> <td>IntraPredAvailFlagA – A (Left Second Half)</td> </tr> </tbody> </table> | Bits | MotionVerticalFieldSelect Index | 7:6 | ChromalIntraPredMode | 5 | Reserved | 4 | IntraPredAvailFlagE – E (Left First Half) | 3 | IntraPredAvailFlagD – D | 2 | IntraPredAvailFlagC – C | 1 | IntraPredAvailFlagB – B | 0 |
| Bits | MotionVerticalFieldSelect Index | | | | | | | | | | | | | | | | |
| 7:6 | ChromalIntraPredMode | | | | | | | | | | | | | | | | |
| 5 | Reserved | | | | | | | | | | | | | | | | |
| 4 | IntraPredAvailFlagE – E (Left First Half) | | | | | | | | | | | | | | | | |
| 3 | IntraPredAvailFlagD – D | | | | | | | | | | | | | | | | |
| 2 | IntraPredAvailFlagC – C | | | | | | | | | | | | | | | | |
| 1 | IntraPredAvailFlagB – B | | | | | | | | | | | | | | | | |
| 0 | IntraPredAvailFlagA – A (Left Second Half) | | | | | | | | | | | | | | | | |
| R+0.4 | 31:16 | LumaIndraPredModes[1] | | | | | | | | | | | | | | | |
| | 15:0 | LumaIndraPredModes[0] | | | | | | | | | | | | | | | |
| R+0.5 | 31:16 | LumaIndraPredModes[3] | | | | | | | | | | | | | | | |
| | 15:0 | LumaIndraPredModes[2] | | | | | | | | | | | | | | | |

(R+0.3 to R+0.5) for an Inter Macroblock

| DWord | Bit | Description |
|-------|-------|------------------------------|
| R+0.3 | 31:24 | Log2WeightDenomChroma |
| | 23:16 | Log2WeightDenomLuma |
| | 15:8 | SubMbPredMode |
| | 7:0 | SubMbShape |
| R+0.4 | 31:16 | LumaIndraPredModes[1] |
| | 15:0 | LumaIndraPredModes[0] |
| R+0.5 | 31:16 | LumaIndraPredModes[3] |
| | 15:0 | LumaIndraPredModes[2] |



The indirect data Motion Vector blocks occupies either 2 or 4 GRF spaces ($n = 2$ or 4) depending on maximum MvSize value. Motion vector blocks are copied from the indirect data buffer without modification, including the case with MvSize = 16 where only Motion Vectors of List 0 is present. VFE is responsible of zero-filling the remainder if the indirect data do not fill up the GRF space n . As noted by (*) in the table below, when $n = 4$, R+2.0 to R+3.7 are zero-filled for MvSize = 0 to 16. If $n = 2$, R+2.0 to R+3.7 are not present.

Table 1-2. AVC-IT indirect data Motion Vector blocks (with $n = 4$)

| DWord | Bit | MvSize = 0 | MvSize = 2 | MvSize = 8 | MvSize = 16 | MvSize = 32 |
|-------|-------|------------|------------|--------------|--------------|--------------|
| R+0.0 | 31:16 | MBZ | MVVert_L0 | MVVert_Y0_L0 | MVVert_Y0_L0 | MVVert_Y0_L0 |
| | 15:0 | MBZ | MVHorz_L0 | MVHorz_Y0_L0 | MVHorz_Y0_L0 | MVHorz_Y0_L0 |
| R+0.1 | 31:16 | MBZ | MVVert_L1 | MVVert_Y0_L1 | MVVert_Y1_L0 | MVVert_Y0_L1 |
| | 15:0 | MBZ | MVHorz_L1 | MVHorz_Y0_L1 | MVHorz_Y1_L0 | MVHorz_Y0_L1 |
| R+0.2 | 31:0 | MBZ | MBZ | MV_Y1_L0 | MV_Y2_L0 | MV_Y1_L0 |
| R+0.3 | 31:0 | MBZ | MBZ | MV_Y1_L1 | MV_Y3_L0 | MV_Y1_L1 |
| R+0.4 | 31:0 | MBZ | MBZ | MV_Y2_L0 | MV_Y4_L0 | MV_Y2_L0 |
| R+0.5 | 31:0 | MBZ | MBZ | MV_Y2_L1 | MV_Y5_L0 | MV_Y2_L1 |
| R+0.6 | 31:0 | MBZ | MBZ | MV_Y3_L0 | MV_Y6_L0 | MV_Y3_L0 |
| R+0.7 | 31:0 | MBZ | MBZ | MV_Y3_L1 | MV_Y7_L0 | MV_Y3_L1 |
| R+1.0 | 31:0 | MBZ | MBZ | MBZ | MV_Y8_L0 | MV_Y4_L0 |
| R+1.1 | 31:0 | MBZ | MBZ | MBZ | MV_Y9_L0 | MV_Y4_L1 |
| R+1.2 | 31:0 | MBZ | MBZ | MBZ | MV_Y10_L0 | MV_Y5_L0 |
| R+1.3 | 31:0 | MBZ | MBZ | MBZ | MV_Y11_L0 | MV_Y5_L1 |
| R+1.4 | 31:0 | MBZ | MBZ | MBZ | MV_Y12_L0 | MV_Y6_L0 |
| R+1.5 | 31:0 | MBZ | MBZ | MBZ | MV_Y13_L0 | MV_Y6_L1 |
| R+1.6 | 31:0 | MBZ | MBZ | MBZ | MV_Y14_L0 | MV_Y7_L0 |
| R+1.7 | 31:0 | MBZ | MBZ | MBZ | MV_Y15_L0 | MV_Y7_L1 |
| R+2.0 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y8_L0 |
| R+2.1 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y8_L1 |
| R+2.2 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y9_L0 |
| R+2.3 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y9_L1 |
| R+2.4 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y10_L0 |
| R+2.5 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y10_L1 |
| R+2.6 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y11_L0 |
| R+2.7 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y11_L1 |
| R+3.0 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y12_L0 |
| R+3.1 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y12_L1 |
| R+3.2 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y13_L0 |



| DWord | Bit | MvSize = 0 | MvSize = 2 | MvSize = 8 | MvSize = 16 | MvSize = 32 |
|-------|------|------------|------------|------------|-------------|-------------|
| R+3.3 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y13_L1 |
| R+3.4 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y14_L0 |
| R+3.5 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y14_L1 |
| R+3.6 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y15_L0 |
| R+3.7 | 31:0 | MBZ (*) | MBZ (*) | MBZ (*) | MBZ (*) | MV_Y15_L1 |

The indirect data Weight-Offset blocks occupies 2 GRF spaces. Data from the indirect data buffer are copied without modification. VFE is responsible of zero-filling the remainder if the indirect data do not fill up the whole 2 GRF space (MvSize is 0 or 2).

AVC-IT indirect data Weight-Offset blocks

| Dword | Bit | MvSize | | |
|-------|-------|--------|-------------|--------------------|
| | | 0 | 2 | 8, 16, 32 |
| R+0.0 | 31:24 | MBZ | Offset_Y_L1 | Offset_Y_Block0_L1 |
| | 23:16 | MBZ | Weight_Y_L1 | Weight_Y_Block0_L1 |
| | 15:8 | MBZ | Offset_Y_L0 | Offset_Y_Block0_L0 |
| | 7:0 | MBZ | Weight_Y_L0 | Weight_Y_Block0_L0 |
| R+0.1 | 31:16 | MBZ | WO_Cb_L1 | WO_Cb_Block0_L1 |
| | 15:0 | MBZ | WO_Cb_L0 | WO_Cb_Block0_L0 |
| R+0.2 | 31:16 | MBZ | WO_Cr_L1 | WO_Cr_Block0_L1 |
| | 15:0 | MBZ | WO_Cr_L0 | WO_Cr_Block0_L0 |
| R+0.3 | 31:0 | MBZ | MBZ | MBZ |
| R+0.4 | 31:16 | MBZ | MBZ | WO_Y_Block1_L1 |
| | 15:0 | MBZ | MBZ | WO_Y_Block1_L0 |
| R+0.5 | 31:16 | MBZ | MBZ | WO_Cb_Block1_L1 |
| | 15:0 | MBZ | MBZ | WO_Cb_Block1_L0 |
| R+0.6 | 31:16 | MBZ | MBZ | WO_Cr_Block1_L1 |
| | 15:0 | MBZ | MBZ | WO_Cr_Block1_L0 |
| R+0.7 | 31:0 | MBZ | MBZ | MBZ |
| R+1.0 | 31:16 | MBZ | MBZ | WO_Y_Block2_L1 |
| | 15:0 | MBZ | MBZ | WO_Y_Block2_L0 |
| R+1.1 | 31:16 | MBZ | MBZ | WO_Cb_Block2_L1 |
| | 15:0 | MBZ | MBZ | WO_Cb_Block2_L0 |
| R+1.2 | 31:16 | MBZ | MBZ | WO_Cr_Block2_L1 |
| | 15:0 | MBZ | MBZ | WO_Cr_Block2_L0 |



| Dword | Bit | MvSize | | |
|-------|-------|--------|-----|-----------------|
| | | 0 | 2 | 8, 16, 32 |
| R+1.3 | 31:0 | MBZ | MBZ | MBZ |
| R+1.4 | 31:16 | MBZ | MBZ | WO_Y_Block3_L1 |
| | 15:0 | MBZ | MBZ | WO_Y_Block3_L0 |
| R+1.5 | 31:16 | MBZ | MBZ | WO_Cb_Block3_L1 |
| | 15:0 | MBZ | MBZ | WO_Cb_Block3_L0 |
| R+1.6 | 31:16 | MBZ | MBZ | WO_Cr_Block3_L1 |
| | 15:0 | MBZ | MBZ | WO_Cr_Block3_L0 |
| R+1.7 | 31:0 | MBZ | MBZ | MBZ |

If MbType != I_PCM, the residual data are fully expanded with zero-fill to take up 24 GRF registers. Each block of 8x8 residual data are stored in GRF in raster-scan order with 16-bit signed integer samples in 2's compliment form.

If MbType = I_PCM, the residual data take up 12 GRF registers. Each block of 8x8 residual data are stored in GRF in raster-scan order with 8-bit unsigned integer samples. The remaining 12 GRF contains undefined values.

AVC-IT Mode indirect data residual data blocks (MbType != I_PCM)

| GRF Registers | Description |
|---------------|--|
| R+0 to R+3 | Block Y0 (Zero fill if skipped) |
| R+4 to R+7 | Block Y1 (Zero fill if skipped) |
| R+8 to R+11 | Block Y2 (Zero fill if skipped) |
| R+12 to R+15 | Block Y3 (Zero fill if skipped) |
| R+16 to R+19 | Block Cb (Zero fill if skipped) |
| R+20 to R+23 | Block Cr (Zero fill if skipped) |

AVC-IT Mode indirect data residual data blocks (MbType == I_PCM)

| GRF Registers | Description |
|---------------|-----------------|
| R+0 to R+1 | Block Y0 |
| R+2 to R+3 | Block Y1 |
| R+4 to R+5 | Block Y2 |
| R+6 to R+7 | Block Y3 |
| R+8 to R+9 | Block Cb |
| R+10 to R+11 | Block Cr |
| R+12 to R+23 | Undefined |



1.8.1.6 AVC-MC Mode Root Thread [DevCTG, DevILK]

The root thread payload messages in AVC-MC mode are identical to that in AVC-IT mode, except the first dword. The first dword is the same as that in AVC-IT mode as shown in **Error! Reference source not found.** See previous section for details for the rest of the parts.

Dword 0 of R+0.0 of AVC-MC In-line data block

| Dword | Bit | Description |
|-------|-------|--|
| R+0.0 | 31:28 | Reserved. MBZ |
| | 26:25 | MbAffFieldFlag This field indicates that the current macroblock is a field macroblock within a MbAff frame picture. It is provided as Flag = MbaffFrame & FieldMbFlag . 00 = if (Flag == 0) 11 = if (Flag == 1) Other encodings are reserved |
| | 24 | FieldMbPolarityFlag This field indicates the field polarity of the current macroblock. Within a MbAff frame picture, this field may be different per macroblock and is set to 1 only for the second macroblock in a MbAff pair if FieldMbFlag is set. Otherwise, it is set to 0. Within a field picture, this field is set to 1 if the current picture is the bottom field picture. Otherwise, it is set to 0. It is a constant for the whole field picture. This field is reserved and MBZ for a progressive frame picture. 0 = Current macroblock is a field macroblock from the top field 1 = Current macroblock is a field macroblock from the bottom field <i>Programming Note: Here bits [26:24] (MbAffFieldFlag and FiedlMbPolarityFlag) match with bits [10:8] of the Media Block Read message descriptor, simplifying the programming for message generation, as when MbAffFieldFlag is "1", kernels need to override the original "frame" surface state set for MBAFF frame picture.</i> |
| | 23 | Reserved. MBZ |
| | 22:17 | Reserved. MBZ (Forced to Zero by VFE) |
| | 16 | Reserved. MBZ |
| | 15 | Transform8x8Flag |
| | 14 | FieldMbFlag (Field Macroblock Flag) This field specifies whether current macroblock is field macroblock. 0 = Frame macroblock. 1 = Field macroblock. |
| | 13 | IntraMacroblock (Intra Macroblock Indicator) |
| | 12:8 | MbType (Macroblock Type) This field, along with "IntraMacroblock" specifies the macroblock types. |
| | 7:6 | WeightedBiPredFlag (Weighted Bidirectional Prediction Flag) (from Picture State). Valid only for macroblock in inter mode. Otherwise (intra macroblock), this field is reserved. |
| | 5 | WeightedPredFlag. (from Picture State). Valid only for macroblock in inter mode. Otherwise (intra macroblock), this field is reserved. |



| Dword | Bit | Description |
|-------|-----|--|
| | 4 | Reserved. MBZ |
| | 3:2 | ChromaFormatIdc (Chroma Format Indicator) This field specifies the chroma format for the decoding process as defined below. This field is constant within a picture. 00 = Luma only (monochrome) 01 = YUV420 sampling 10 is reserved (for YUV422 sampling) 11 is reserved (for YUV444 sampling) |
| | 1 | MbaffFrame. (from Picture State). |
| | 0 | FieldPicFlag. (from Picture State) |

1.8.1.7 VC1-IT Mode Root Thread [DevCTG,, DevILK]

The following table shows the root thread payload messages when VFE is in VC1-IT mode. One example GRF register location is given for the condition that: CURBE is disabled.

VC1-IT Mode root thread payload layout

| GRF Register | Example | Description |
|------------------|----------|--|
| R0 | R0 | R0 header |
| R1 – R(m) | n/a | Constants from CURBE when CURBE is enabled m is a non-negative value. If m is 0, this field doesn't exist. |
| R(m+1) | R1 | In-line Data block. |
| R(m+2) – R(m+13) | R2 – R13 | Indirect Data Residual blocks |

The R0 header field is the same as in other modes.

The in-line data block field is the same as in the MEDIA_OBJECT_EX command.

1.8.1.8 Child Thread

The thread initiation for the child thread is determined by the data stored in the URB by the parent that spawns it. No hardware-defined header is generated. Software should follow the header field definition similar to that for a root thread, when the same fields are used, to be consistent and to reduce message header assemble overhead.

The Parent Thread Count field should be the Thread Count field of the parent thread itself (e.g. copying R0.6[23:0] to R0.7[23:0]). The Thread Count field should have a unique value for each child thread and the unique value should not be dependent on the execution order. This is mostly important for the cases when the child thread generation order may vary depending on the thread completion order. For example, when generating child threads for macroblock-based processing, the Thread Count field for a child thread should be deterministic for a macroblock position.



The following table shows the R0 register contents for a child thread, which is generated by its parent thread. The remaining payloads are application dependent.

| DWord | Bit | Description |
|-----------|-------|-------------------------|
| R0.7 | 31 | Reserved |
| | 27:24 | Reserved |
| | 23:0 | Reserved |
| R0.6 | 31:24 | Reserved |
| | 23:0 | Reserved |
| R0.5-R0.0 | 31:0 | Software defined |

1.8.1.9 GPGPU Thread

The RO header of the Thread Dispatch Payload for the GPGPU thread:

| DWord | Bit | Description |
|-------|-------|---|
| R0.7 | 31:0 | Thread Group ID Z: This field identifies the Z component of the thread group. That this thread belongs to. |
| R0.6 | 31:0 | Thread Group ID Y: This field identifies the Y component of the thread group that this thread belongs to. |
| R0.5 | 31:10 | Scratch Space Pointer. Specifies the 1k-byte aligned pointer to the scratch space (used for the GPGPU local memory space). Format = GeneralStateOffset[31:10] |
| | 9:8 | Reserved : MBZ |
| | 7:0 | FFTID. This ID is assigned by TS and is a unique identifier for the thread in comparison to other concurrent threads (of any thread group). It is used to free up resources used by the thread upon thread completion. |
| R0.4 | 31:5 | Binding Table Pointer: Specifies the 32-byte aligned pointer to the Binding Table. It is specified as an offset from the Surface State Base Address . Format = SurfaceStateOffset[31:5] |
| | 4:0 | Reserved : MBZ |
| R0.3 | 31:5 | Sampler State Pointer. Specifies the 32-byte aligned pointer to the sampler state table. Format = GeneralStateOffset[31:5] |
| | 4 | Reserved : MBZ |
| | 3:0 | Per Thread Scratch Space. Specifies the amount of scratch space, in 16-byte quantities, allowed to be used by this thread. The value specifies the power that two will be raised to, to determine the amount of scratch space. Format = U4 Range = [0,11] indicating [1k bytes, 2M bytes] in powers of two |
| R0.2 | 31:28 | Reserved: MBZ |



| DWord | Bit | Description |
|-------|-------|--|
| | 27:24 | BarrierID: This field indicates the barrier that this kernel is associated with. Format: U4 |
| | 23:16 | Reserved: MBZ |
| | 15:9 | Reserved : MBZ |
| | 8:4 | Interface Descriptor Offset. This field specifies the offset from the interface descriptor base pointer to the interface descriptor which will be applied to this object. It is specified in units of interface descriptors. Format = U5 |
| | 3:0 | Reserved. MBZ |
| R0.1 | 31:0 | Thread Group ID X: This field identifies the X component of the thread group that this thread belongs to. |
| R0.0 | 31:28 | Reserved: MBZ |
| | 27:24 | Shared Local Memory Index: Indicates the starting index for the shared local memory for the thread group. Each index points to the start of a 4k memory block, 16 possibilities cover the entire 64k shared memory per half-slice. Format = U4 |
| | 23:16 | Reserved: MBZ |
| | 15:0 | URB Handle. This is the URB handle where indicating the URB space for use by the thread. |

R1 and above contain the X/Y/Z values for each channel in the thread. The exact format is determined by the compiler.

1.8.2 Thread Spawn Message

The thread spawn message is issued to the TS unit by a thread running on an EU. This message contains only one 8-DW register. The thread spawn message may be used to

- Spawn a child thread
- Spawn a root thread (start dispatching a synchronized root thread)
- Dereference URB handle
- Indicate a thread termination, dereference other TS managed resource and may or may not dereference URB handle

In order to end a root thread, the end of thread message must be targeted at the thread spawner. In this case, the root thread sends a message with a “dereference resource” in the Opcode field. The thread spawner does *not* snoop the messages sideband to determine when a root thread has ended. Thread Spawner does not track when a child thread terminates, to be consistent a child thread should also terminate with a “dereference resource” message to the Thread Spawner. Software must set the Requester Type (root or child thread) field correctly.

As TS dispatches one synchronized root thread upon receiving a ‘spawn root thread’ message (from a synchronization thread). The synchronizing thread must send the number of ‘spawn root thread’ message exactly the same as the subsequent ‘synchronized root thread’. No more, no less. Otherwise, hardware behavior is undefined.



URB Handle Offset field in this message (in M0.4) has 10 bits, allowing addressing of a large URB space. However, when a parent thread writes into the URB, it subjects to the maximum URB offset limitation of the URB write message, which is only 6 bits (see Unified Return Buffer Chapter for details). In this case, the parent thread may have to modify the URB Return Handle 0 field of the URB write message in order to subdivide the large URB space that the thread manages.

[DevILK]: In addition to monitor ‘end of thread message’ targeted to Thread Spawner, Thread Spawner also monitors the message targeting to Message Gateway for EOT signal. Therefore, a child thread, who doesn’t hold any hardware resource (URB handle or scratch memory) that Thread Spawner manages, can terminate with a Gateway message with EOT on. The reason of this new TS feature is to avoid a possible risk condition as described below.

[DevILK]: In addition to monitor ‘end of thread message’ targeted to Thread Spawner, Thread Spawner also monitors the message targeting to Message Gateway for EOT signal. Therefore, a child thread, who doesn’t hold any hardware resource (URB handle or scratch memory) that Thread Spawner manages, can terminate with a Gateway message with EOT on. The reason of this new TS feature is to avoid a possible risk condition as described below.

In a system running child threads, a parent thread is monitoring the status of the child threads by communications through Message Gateway. When a child thread is about to terminate, it sends a message to the parent through Message Gateway and then sends a second message of EOT (end of thread) to TS.

There is a latency between sending a message to parent thread and the EOT to TS due to message bus arbitration. The parent thread may acknowledge the GW message and issue a new child dispatch before the EOT was processed; basically threads are issued faster than retired.

Because the messages for new child dispatch and EOT go to the same queue in TS, if the queue gets full, EOTs will get blocked. In the case when all the EUs/Threads are full, this will create a system deadlock: no EOTs can be acknowledged by TS (to free up EU resource) and no child threads can be dispatched (to free up TS queue to receive EOT message).

1.8.2.1 Message Descriptor

The following table shows the lower 16 bits of the message descriptor (lower 20 bits for [DevILK+]) within the SEND instruction for a thread spawn message.

| Bit | Description |
|------|---|
| 19 | <p>[DevILK+]: Header Present This bit must be set to zero for all Thread Spawner messages.</p> <p>[Pre-DevILK]: this bit is not part of the shared function specific message descriptor.</p> |
| 18:5 | <p>Reserved: MBZ</p> <p>[Pre-DevILK]: Bits 18:16 are not part of the shared function specific message descriptor.</p> |
| 4 | <p>Resource Select. This field specifies the resource associated with the action taken by the Opcode.</p> <p>If Opcode is “Spawn thread”, this field selects whether it is a child thread or a root thread.</p> <p>0: spawn a <i>child</i> thread</p> <p>1: spawn a <i>root</i> thread or ([DevCTG] only) release a PRT_Fence</p> <p>If Opcode == “Dereference Resource”, this field indicates whether the URB handle is to be dereferenced. The URB handle can only be dereferenced once.</p> <p>0: The URB handle is dereferenced</p> |



| | |
|-----|--|
| | 1: The URB handle is NOT dereferenced |
| 3:2 | Reserved: MBZ |
| 1 | <p>Requester Type. This field indicates whether the requesting thread is a root thread or a child thread. If it is a root thread, when Opcode is 0, FF managed resources will be dereferenced. If it is a child thread and Opcode is 0, no resource will be dereferenced – basically no action is required by the TS.</p> <p>0: Root thread 1: Child thread</p> |
| 0 | <p>Opcode. Indicates the operation performed by the message. A root thread must terminate with a message to TS (Opcode == 0 and EOT == 1). A child thread <i>should</i> also terminate with such a message. A thread cannot terminate with an Opcode of "spawn thread".</p> <p>0: dereference resource (also used for end of thread) 1: spawn thread</p> |

1.8.2.2 Message Payload

| DWord | Bit | Description |
|-------|-------|---|
| M0.7 | 31:0 | Reserved |
| M0.6 | 31:0 | Reserved |
| M0.5 | 31:8 | Ignored |
| | 7:0 | <p>FTID. This ID is assigned by TS and is a unique identifier for the thread in comparison to other concurrent root threads. It is used to free up resources used by a root thread upon thread completion.</p> <p>This field is valid only if the Opcode is "dereference resource", and is ignored by hardware otherwise.</p> |
| M0.4 | 31:16 | Ignored |
| | 15:10 | <p>Dispatch URB Length. Indicates the number of 8-DW URB entries contained in the Dispatch URB Handle that will be dispatched. When spawning a child thread, the URB handle contains most of the child thread's payload including R0 header. When spawning a root thread, the URB handle contains the message passed from the requesting thread to the spawned "peer" root thread. The number of GRF registers that will be initialized at the start of the spawned child thread is the addition of this field and the number of URB constants if present. The number of GRF registers that will be initialized at the start of a spawned root thread is the addition of this field, the number of URB constants if present, and the URB handle received from VFE.</p> <p>This field is ignored if the Opcode is "dereference resource".</p> <p>Length of 0 can be used while spawning child threads to indicate that there is no payload beyond the required R0 header. Length of 0 while spawning a root thread indicates that there is no payload at all from the parent thread. A spawned root has R0 supplied by the Media_Object command indirect/inline data.</p> <p>Format = U6 Range = [0,63] for child threads Range = [1,63] for root threads</p> |



| DWord | Bit | Description |
|-------------------|-------|--|
| | 9:0 | <p>URB Handle Offset. Specifies the 8-DW URB entry offset into the URB handle that determines where the associated dispatch payload will be retrieved from when the spawned child or root thread is dispatched.</p> <p>This field is ignored if the Opcode is “dereference resource”.</p> <p>Format = U10</p> <p>Range = [0,1023]</p> |
| M0.3 | 31:0 | Ignored |
| M0.2 [Pre-DevSNB] | 31:4 | <p>Interface Descriptor Pointer. Specifies the 16-byte aligned pointer to <i>the child thread’s</i> interface descriptor. This pointer is used by TS to fetch the interface descriptor for the child thread, and it is also passed to the child thread in its R0 header.</p> <p>This field is ignored if the Opcode is “dereference resource” or “spawn a root thread”.</p> <p>Format = GeneralStateOffset[31:4]</p> |
| | 3:0 | Ignored |
| | 27:24 | <p>BarrierID. This field indicates which one from the 16 Barriers this kernel is associated.</p> <p>Format: U4</p> |
| | 23:16 | <p>Barrier.Offset. This is the offset for the Barrier to indicate the offset from the requester’s RegBase (which may be 0 if Bypass Gateway Control is set to 1) for the broadcast barrier message. Barrier.Offset + RegBase must be in the valid GRF range. Otherwise, hardware behavior is undefined.</p> <p>It is in unit of 256-bit GRF register.</p> <p>The most significant bit of this field must be zero.</p> <p>Format = U8</p> <p>Range = [0,127]</p> |
| | 15:9 | Ignored |
| | 8:4 | <p>Interface Descriptor Offset. This field specifies the offset from the interface descriptor base pointer to the interface descriptor which will be applied to this object. It is specified in units of interface descriptors.</p> <p>Format = U5</p> |
| | 3:0 | <p>Scoreboard Color (only with MEDIA_OBJECT_EX): This field specifies which dependency color the current thread belongs to. It affects the dependency scoreboard control.</p> <p>Format = U4</p> |
| M0.1 | 31:0 | Ignored |
| M0.0 | 31:28 | Ignored |
| | 27:24 | <p>[Ivb+]Shared Local Memory Index: Indicates the starting index for the shared local memory for the thread group. Each index points to the start of a 4k memory block, 16 possibilities cover the entire 64k shared memory per half-slice.</p> <p>Format = U4</p> <p>[pre-Ivb] Reserved: MBZ</p> |
| | 23:16 | <p>[DevIL+] Thread Dependency Identifier (TDID). This field is assigned by TS to be used for hardware scoreboard.</p> <p>Format = U8</p> <p>[Pre-DevIL] Reserved : MBZ</p> |



| DWord | Bit | Description |
|-------|------|---|
| | 15:0 | <p>Dispatch URB Handle</p> <p>If Opcode (and Requester Type) is "spawn a child thread": Specifies the URB handle for the child thread.</p> <p>If Opcode (and Requester Type) is "spawn a root thread": Specifies the URB handle containing message (e.g. requester's gateway information) from the requesting thread to the spawned root thread.</p> <p>If Opcode is "dereference resource": This field is required on end of thread messages if the Children Present bit is set, as the handle must be dereferenced, otherwise this field is ignored.</p> |



2. AVC Bit-Serial Decoder [DevCTG/DevILK]

2.1 Design Assumptions

MbaffFrameFlag can't have different values in different slices of the same picture

512 bits memory access interface, cache line size

No interruptability within a slice decoding. In addition, according to the AVC specification, there is no dependency across Slice boundary with respect to bit stream decoding. Hence, context switching mechanism is much simplified. Currently, a video stream context is consisted of mainly the starting MB number of a Slice in that stream. Context switching mechanism is needed for handling concurrent multiple stream decoding and error handling.

There is no pre-emption for the BCS pipeline (CPU is not able to interrupt the BCS-BSD pipe), hence every command buffer is required to contain all the states setup (preamble), and therefore no need to save and restore context.

If different H.264 applications are running, each will have its own command buffer (responsible for decoding at least one Slice). The driver is required to collect all the command buffers of a picture of a video sequence. Each command buffer will run to completion serially.

If an H.264 application is playbacking multiple video sequences, each sequence has its own command buffer list.

The switching across different video streams of different applications can only be done on a picture boundary.

A command buffer can contain multiple Slices to be decoded, but will not cross the picture boundary. But there is no driver interruption in between Slices. A command buffer must run to its completion. The worst case latency is one picture time, since a Slice can be as big as a picture.

2.2 Bit-Serial Decoding (BSD) Unit

A standalone H/W unit uses for performing the AVC CABAC/CAVLD decoding function. It is operating concurrently with the GEN4 GPE, with its own Command Streamer. The communications and synchronization between BSD and GEN4 GPE is primarily through memory. The Host will parse all header information from an AVC bit-stream, and only send down the raw Slice Data in memory to the BSD. The decoded data are written back to memory for the GEN4 GPE to consume.

Starting code detection is done in the Host driver, only the Slice Data layer is passed down to the HW for processing.

The BSD gets its commands, states and parameters from the CS unit. AI unit now fetches/pre-fetches the indirect slice payload through AM/CS/CI units and takes care of start/end of slice before sending the indirect data (slice payload) to AC unit. AI unit also needs to take care of emulation prevention byte for AC unit.



2.2.1.1 4x4 Spatial Luma, 16x16 Spatial Luma, 8x8 Spatial Chroma

Intra_16x16 uses one intra prediction scheme for the whole macroblock. Pixels may be filled from surrounding macroblocks at the left and the upper edge using one of four possible prediction modes. Intra prediction is also performed for the chroma planes using the same range of prediction modes. However, different modes may be selected for luma and chroma.

Intra_4x4 subdivides the macroblock into 16 subblocks and assigns one of nine prediction modes to each of these 4x4 blocks. The prediction modes offered in Intra_4x4 blocks support gradients or other smooth structures that run in one of eight distinct directions. One additional mode fills a whole subblock with a single value and is used if no other mode fits the actual pixel data inside a block. Intra chroma prediction is performed in an identical way to the one used in Intra_16x16 prediction. Thus, the chroma predictions are not split into subblocks as it is done for luma.

2.2.1.2 PCM

Transmit picture samples (Luma and Chroma) directly without any prediction, transformation and compression. It sets the upper limit (maximum number of bits) for each coded MB. The I_PCM mode is a type of "fail safe" mode that bounds the size of a macroblock in the compressed video bitstream. In I_PCM, coefficient data is replaced by actual 8-bit pixel sample values. Hence, instead of 16-bit per coefficient data, a MB in I_PCM mode will contain packed 8-bit values of Luma and Chroma, and each MB will be exactly 384 bytes (luma + chroma) in size.

According to the AVC Spec., I_PCM mode is enlisted amongst Intra-Prediction modes. And internal parameter settings (listed below) in I_PCM mode are similar to those of Intra16x16 mode.

For a macroblock coded with mb_type equal to I_PCM, the Intra macroblock prediction mode shall be inferred.

Hence, there will be no separation of DC terms from the AC terms as in inter-prediction coding

So, all the Dcblock flags should be either ignored by the H/W in the I_PCM mode (as in intra-prediction modes) or set to 0 (to indicate their absence)

For macroblocks in I_PCM mode, there is no coded_block_pattern syntax element present in the bitstream to derive the variables CodedBlockPatternLuma and CodedBlockPatternChroma.

According to Table 7-11, when in I_PCM mode, the parameters - Intra16x16PredMode, CodedBlockPatternChroma and CodeBlockPattern Luma are don't care.



Hence, in I_PCM mode, we treat each MB as non-skipped and all 16x16 Luma coefficients and 8x8 Chroma (Cr and Cb) coefficients are present and no separation of DC terms. So, the cbp bits for all Luma and Cr/Cb 8x8 blocks should be set to indicate their presence. In monochrome mode, Cr/Cb should be set to absent. And the separated DC block flags should be 0 to indicate their absence.

When `chroma_format_idc` is equal to 0 (i.e., monochrome), `CodedBlockPatternChroma` shall be equal to 0.

| | |
|---|--|
| <code>DcBlockCodedYFlag</code> [Used by VFE] | Always zero? |
| <code>DcBlockCodedCbFlag</code> [Used by VFE] | Always zero? |
| <code>DcBlockCodedCrFlag</code> [Used by VFE] | Always zero? |
| <code>CbpY</code> (Coded Block Pattern Y) | Always 0xFF? |
| <code>CbpCr</code> [Used by VFE] | Always 0xF in non-monochrome mode, 0x0 in monochrome mode? |
| <code>CbpCb</code> [Used by VFE] | Always 0xF in non-monochrome mode, 0x0 in monochrome mode? |

2.3 H/W Asynchronous Reset

Global H/W Reset is signalled during system power-up to reset the entire BSD pipeline and to initialize it to a known state. It is also used in situation of errors to clear the pipeline. When the Command Stream Parser is stalled, pipeline reset implemented as CS command would not get through. Hence, the asynchronous H/W Reset implemented through MMIO register and issued by the Host S/W is the only option to get out. Host S/W WatchDog timers may be implemented to detect stalled conditions. After H/W reset, the ring buffer will be lost, and driver re-programming is required to continue. The reset immediately takes effect regardless the current operation.

2.4 H/W Pipeline Flush

A flush is implicitly carried out at the end of decoding a Slice. An explicit flush can also be issued to the BCS at the picture boundary (at the end of processing a picture). BCS unit will then perform the flush by monitoring the “done” signals from AI, AC and AM units, and issue a subsequent “store Dword”.

2.5 MMIO Interface

MMIO is involved in interrupt generation during error handling. MMIO address 4400h is defined for the BSD registers, and 4000h for the BCS MMIO. These are read and write registers.

Currently, there are two set of MMIO registers

- AVC BSD Error Register (16-bit)
- AVC BSD Error Count Register (12-bit each) – do not wrap around when maximum count has reached.

AVC BSD Error Register : MMIO Address : 4400h

The only MMIO write operation is to reset each individual bit of this register to a 0 value, after the error information has been read and/or processed. The driver may choose to read this register in between pictures and video sequence and upon video stream switching. This register is set to 0 at powerup.



16-bit Read/Write Register (32-bit aligned)

| Bit | Description |
|------|---|
| 15:4 | Reserved. MBZ |
| 3 | BSD Premature Completion Error Status Flag When a BSD Premature Completion error has occurred and the BSDPrematureComplete Error Handling bit in the inline data of the AVC_BSD_OBJECT command is set, this error status flag is set until being cleared by a subsequent MMIO write to this register. |
| 2 | MPR Error Status Flag When a MPR error has occurred and the MPR Error Handling bit in the inline data of the AVC_BSD_OBJECT command is set, this error status flag is set until being cleared by a subsequent MMIO write to this register. |
| 1 | VLD Error Status Flag When a VLD error has occurred and the VLD Error Handling bit in the inline data of the AVC_BSD_OBJECT command is set, this error status flag is set until being cleared by a subsequent MMIO write to this register. |
| 0 | BSD Error Status Flag When a BSD error has occurred and the BSD Error Handling bit in the inline data of the AVC_BSD_OBJECT command is set, this error status flag is set until being cleared by a subsequent MMIO write to this register. |

AVC BSD Error Count Register : MMIO Address : 4404h

The only MMIO write operation is to reset this register to a 0 value. The driver may choose to read this register in between pictures and video sequence and upon video stream switching. This register is set to 0 at powerup.

16-bit Read/Write Register (32-bit aligned)

| Bit | Description |
|-------|---|
| 15:12 | Reserved. MBZ |
| 11:0 | BSD Error Count Increment by 1 when any of the recognized errors (BSD, VLD, MPR and PrematureCompletion) has occurred. Do not wrap around when the maximum count has reached. |

2.6 Programming the BSD Unit

2.6.1 Example Command Sequence for Decoding a Single Video Stream

BSD Unit uses a simple programming model that does not support pipelined state changes. All state information required to decode a Slice are always sent along (and prior to) with the compressed video data. This eliminates the need to reset the internal state registers for decoding a new Slice; they will be over-written anyway.



The basic steps in programming the BSD Unit are listed below. There is no explicit initialization step for the BSD engine, the H/W unit will be properly initialized itself at power-up and at reset. Some of the steps are optional and some are repetitive. The recommended order should be followed strictly. Some usage restrictions are highlighted for illustration purpose. For details, reader should refer to the respective chapters for these commands.

- Step 1: Issue the AVC_IND_OBJ_BASE_ADDR command for system configuration
 - This command is mandatory for this step (i.e. at least one).
 - Multiple such commands in this step are allowed. The last one overwrites previous ones.
 - This command must precede any other state commands below.
 - The fields **AVC Indirect Object Base Address** and **AVC Indirect Object Access Upper Bound** are used to control indirect object load.
- Step 2: Issue AVC_BSD_State commands to program the AVC decoding pipeline states
 - These commands are mandatory, since Slice Data decoding cannot rely on states set by the previous Slice.
 - Multiple such commands in this step are allowed. When the same command is issued multiple times, the last one overwrites previous ones.
 - There are 4 state commands to be sent in a strict order, Except the IMG_STATE that should only be changed on a per picture basis, the others are re-sent for decoding a new Slice whenever there is a change :
 - AVC_BSD_IMG_STATE – sent in for every new picture (first slice of each picture)
 - AVC_BSD_QM_STATE – sent in for every Slice to be decoded
 - AVC_BSD_SLICE_STATE – sent in for every Slice to be decoded
 - AVC_BSD_BUF_BASE_STATE – sent in every Slice to be decoded
- Step 3: Issue the media/primitive AVC_BSD_OBJECT command
 - Multiple AVC_BSD_OBJECT commands can be issued to continue processing BSD media primitives, only if all the state information is unchanged.
 - An implicit pipeline flush is executed at the end of decoding a Slice. The BSD Command Stream Parser will be halted until it has received a complete signal from the BSD unit.
- Step 4: Repeat Step 2 and 3 as many times as needed for processing an entire picture
- Step 5: Send a phantom slice at the end of each picture
 - To make sure that the automatic error concealment in the H/W is triggered to the end of the picture when needed.
- Step 6 : BCS flush sent in at the end of each picture, and perform a subsequent store Dword command
- Step 7: Repeat Step 2 through 5 as many times as needed for processing an entire video sequence



2.7 AVC_BSD Commands

2.7.1 BSD_IND_OBJ_BASE_ADDR Command

The BSD_IND_OBJ_BASE_ADDR command sets the memory base address pointers for the subsequent Indirect Data Start Address specified in the AVC_BSD_OBJECT commands. This command is shared by VC1 BSD pipe.

While the use of this base addresses is unconditional, the indirection can be effectively disabled by setting the base addresses to zero.

The Command Streamer (BCS) will perform the memory access bound check automatically using the AVC Indirect Object Access Upper Bound specification. If any access is at or beyond this bound, zero value is returned. The request to memory still being sent, but the BSD H/W will detect and perform the zeroing. If the Upper Bound is turned off, the beyond bound request will return whatever on the bus (invalid data).

| Notation | Definition |
|----------------------|---|
| PhysicalAddress[n:m] | Corresponding bits of a physical graphics memory byte address (not mapped by a GTT) |
| GraphicsAddress[n:m] | Corresponding bits of an absolute, virtual graphics memory byte address (mapped by a GTT) |

| Dword | Bit | Description |
|-------|-------|--|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | AVC Command Opcode = BSD_IND_OBJ_BASE_ADDR Pipeline[28:27] = BSD = 2h; Opcode[26:24] = AVC = 4h; Sub Opcode[23:16] = 4h |
| | 15:0 | DWord Length (Excludes DWords 0, 1) = 0001h |
| 1 | 31:12 | AVC Indirect Object Base Address. Specifies the 4K-byte aligned memory base address for indirect object load in AVC_BSD_OBJECT command. Format = GraphicsAddress[31:12] |
| | 11:0 | Reserved : MBZ |
| 2 | 31:12 | AVC Indirect Object Access Upper Bound. This field specifies the 4K-byte aligned (exclusive) maximum Graphics Memory address access by an indirect object load in a AVC_BSD_OBJECT command. Indirect data accessed at this address and beyond will appear to be 0. Setting this field to 0 will cause this range check to be ignored. If non-zero, this address must be greater than the AVC Indirect Object Base Address . Hardware ignores this field if indirect data is not present. Format = GraphicsAddress[31:12] |
| | 11:0 | Reserved: MBZ |



2.7.2 AVC_BSD_STATE Commands

The AVC_BSD_STATE commands are used to load various state information and parameters into the BSD unit to govern its operations. They include the decoding parameters extracted from different AVC syntax layers (e.g. NAL, PPS, SPS, Slice Header, SEI, etc.) that are to be consumed within the BSD (and MPR) unit, as well as decoding information that are used to drive and to be passed along to other processing stages (IT, MC and ILDB) that follow the BSD unit. Some of these state parameters are directly set to the corresponding AVC syntax elements' values; and others (with similar name as the corresponding AVC syntax element) are derived from them as indicated. There are also additional decoding parameters that are not defined as part of the AVC specification, but are needed for proper functioning (e.g. error handling state parameters).

Each state command is of different size, and is designed to be Dword (32-bit) aligned. Some of these parameters are defined in according to the AVC specification, and some are only related to the current implementation of the BSD unit and Intel architecture.

These commands are issued prior to a set of BSD Object commands, so that all the required state variables are set appropriately prior to the actual decoding of the corresponding Slice Data portion of the bitstream. The BSD unit is a stateless machine, all state information and parameters must be set prior to the decoding. However, unless a reset is triggered, the current parameter values (programmed from a previous state command) remain until another corresponding AVC_BSD_STATE command comes in to re-program them.

There is an inherent order of setting the state variables, as dedicated by the H/W interface, and therefore the order of issuing the corresponding AVC_BSD_STATE commands. In general, we will follow the order listed below:

1. AVC_BSD_IMG_STATE (fixed size)
2. AVC_BSD_QM_STATE (variable size)
3. AVC_BSD_SLICE_STATE (variable size with fixed sized data structures)
4. AVC_BSD_BUF_BASE_STATE (fixed size)

All state information are sent to the BSD unit through AVC_BSD_STATE commands as inline data (some are fixed size, and some are variable in length but are predetermined by the driver before issuing the corresponding STATE Command).

The current active SPS can be chosen from an array of SPS with an index in the driver. Likewise, the current active PPS can be chosen from an array of PPS with an index in the driver.

2.7.2.1 AVC_BSD_IMG_STATE Command

AVC_BSD_IMG_STATE command encapsulates the decoding parameters that are read or derived from bitstream syntax elements of a NAL unit, a PPS unit, a SPS unit and a Slice Header (only includes those that are invariant within a picture). It includes all the relevant settings for the Frame Parameter and the Image Parameter Interface of the BSD Unit. These parameters are not changes from slice to slice of the same picture, but may change from picture to picture. Hence, this command is only issued at the beginning of processing a new picture and prior to the AVC_BSD_OBJECT command. It has a fixed size of 8 DWs (including the Command Opcode DW0).

AVC_BSD_IMG_STATE command must be sent in order for the H/W to sync to the beginning of a picture.

The values set for these state variables are retained internally across Slices, until they are reset by H/W Asynchronous Reset or changed by the next AVC_BSD_IMG_STATE command.



| Dword | Bit | Description |
|-------|-------|--|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | AVC Command Opcode = AVC_BSD_IMG_STATE Pipeline[28:27] = BSD = 2h; Opcode[26:24] = AVC = 4h; Sub Opcode[23:16] = 0h |
| | 15:0 | DWord Length (Excludes DWords 0,1) = 0004h |
| 1 | 31:16 | Reserved. MBZ. |
| | 15:0 | Frame size in MB unit, FrameSizeInMBs[15:0] Note: bit 15 must be 0 and is ignored by the current H/W implementation, i.e. only bits[14:0] are significant. Unsigned integer value. The value for FrameSizeInMBs must match the product of FrameWidthInMBs and FrameHeightInMBs. Max. screen resolution is therefore limited to 2K x 2K in MB unit. e.g for 1920x1080, FrameSizeInMBs[15:0] = 8160 (1920/16 * 1088/16; rounded up 1080) This parameter is specified for Intel interface only. |
| 2 | 31:24 | Reserved. MBZ. |
| | 23:16 | Frame height in MB unit, FrameHeightInMBs[7:0] Unsigned integer value. It is set to the value of (FrameHeightInMBsMinus1+ 1). Since the max value for FrameHeightInMBs is 255, the max allowed value for FrameHeightInMBsMinus1 is only 254. The min value for FrameHeightInMBs is 1. Although the max. value that can be specified for FrameHeightInMBs is 255 (in the current implementation), FrameWidthInMBs * FrameHeightInMBs must not exceed the max value of FrameSizeInMBs[14:0]. e.g., for 1920x1080, FrameHeightInMBs[7:0] is equal to 68 (1080 divided by 16, and rounded up, i.e. effectively specified as 1088 instead). It is derived from FrameHeightInMbs = (2 – frame_mbs_only_flag) * PicHeightInMapUnits (firmware ???), or PicHeightInMbs = FrameHeightInMbs / (1 + field_pic_flag) internally done (???) |
| | 15:8 | Reserved. MBZ. |
| | 7:0 | Frame width in MB unit, FrameWidthInMBs[7:0] Unsigned integer value. It is set to the value of (FrameWidthInMBsMinus1 + 1). FrameWidthInMBsMinus1 is equal to the value of the syntax element in the current active SPS. Since the max value for FrameWidthInMBs is 255, the max value allowed for FrameWidthInMBsMinus1 is only 254. The min value for FrameWidthInMBs is 1. Although the max. value that can be specified for FrameWidthInMBs is 255 (in the current implementation), FrameWidthInMBs * FrameHeightInMBs must not exceed the max value of FrameSizeInMBs[14:0]. e.g., for 1920x1080, FrameWidthInMBs[7:0] is equal to 120 (1920 divided by 16). |
| 3 | 31:29 | Reserved. MBZ |



| Dword | Bit | Description |
|-------|-------|---|
| | 28:24 | <p>Second Chroma QP Offset, second_chroma_qp_offset[4:0] Signed integer value. It should be in the range of -12 to +12 (according to AVC spec). It specifies the offset for determining QP Cr from QP Y. It is set to the upper 5 bits of the value of the syntax element (Chroma_qp_offset[9:0]) read from the current active PPS. Chroma_qp_offset [4:0] – chroma_qp_offset_bits (from the current active PPS) Chroma_qp_offset [9:5] – second_chroma_qp_offset_bits</p> |
| | 32:21 | Reserved. MBZ |
| | 20:16 | <p>Chroma QP Offset, chroma_qp_offset[4:0] Signed integer value. It should be in the range of -12 to +12 (according to AVC spec). It specifies the offset for determining QP Cb from QP Y. It is set to the lower 5 bits of the value of the syntax element (Chroma_qp_offset[9:0]) read from the current active PPS. Chroma_qp_offset [4:0] – chroma_qp_offset_bits (from the current active PPS) Chroma_qp_offset [9:5] – second_chroma_qp_offset_bits</p> |
| | 15 | <p>Special Kernel MB Scan Order for AVC ILDB Data Writes It controls the way of writing the generated ILDB Data package to memory. 0 – MB are arranged in Raster Scan Order 1 – MB are arranged in the Special Scan Order This parameter is specified for Intel interface only.</p> |
| | 14 | <p>Special Kernel MB Scan Order for AVC-IT Command Writes It controls the way of writing the generated AVC-IT Command package to memory. 0 – MB are arranged in Raster Scan Order 1 – MB are arranged in the Special Scan Order This parameter is specified for Intel interface only.</p> |
| | 13 | <p>Special Kernel MB Scan Order for AVC-IT Data Writes It controls the way of writing the generated AVC-IT Data package to memory. 0 – MB are arranged in Raster Scan Order 1 – MB are arranged in the Special Scan Order This parameter is specified for Intel interface only.</p> |
| | 12 | <p>Monochrome Prediction Weighting Table Decoding Mode Flag, monochrome_pwt_flag Added to cover unapproved H.264 Spec. modification. Should always be written as "1". This parameter is specified for Intel interface only.</p> |
| | 11 | Reserved. MBZ |



| Dword | Bit | Description |
|-------|-----|---|
| | 10 | <p>QM Present Flag, qm_present_flag</p> <p>0 – no Scaling Matrix is present nor sent to the BSD unit; use the flat4x4 (when transform_8x8_mode_flag == 0) or flat8x8 (when transform_8x8_mode_flag == 1) built-in QM matrices. That is, no QM_State command should follow, or the BSD unit should ignore it if it is issued (the BCS will continue to parse QM_State command, but the BSD unit will ignore and not to receive the incoming QM matrices data).</p> <p>1 – Scaling Matrix may be present, depending on the use_default and list_present flags present in the inline data of a subsequent QM_STATE command. A QM_STATE command must follow to complete the specification of the QM matrices for the current picture decoding.</p> <p>It is derived from the pic_scaling_matrix_present_flag in the current active PPS and the seq_scaling_matrix_present_flag in the current active SPS.</p> <p>This parameter is specified for Intel interface only.</p> |
| | 9:8 | <p>Image Structure, img_structure[1:0]</p> <p>The current decoding picture structure can only take on 3 possible values :</p> <ul style="list-style-type: none"> 00 – Frame Picture 01 – Top Field Picture 11 – Bottom Field Picture 10 – Invalid, not allowed. <p>img_structure[0] can be used as a flag to distinguish between frame and field structure. It must be consistent with the field_pic_flag setting in the Slice Header.</p> <p>This parameter is specified for Intel interface only.</p> |



| Dword | Bit | Description |
|-------|-------|--|
| | 7:0 | <p>Current Decoded Image Frame Store ID, <code>img_dec_fs_idc</code> Unsigned integer value.</p> <p>It specifies the destination (currently decoded) picture by its binding table index (Intel implementation) of the current DPB.</p> <p>This parameter is programmed in frame. <code>img_dec_fs_idc[4:0]</code> can only be equal to 0 to 16. Bit[7:5] must be 0.</p> <p>When it is in the range of 0 to 15, hardware uses the selected address from <code>direct_mv_rd0</code> to <code>direct_mv_rd31</code> (based on <code>img_structure</code> and <code>img_dec_fs_idc</code>) as the Direct MV Write Base Address for the Current Picture.</p> <p>When it is 16, hardware uses the <code>direct_mv_wr0_top</code> or <code>direct_mv_wr0_bottom</code> (depending on <code>img_structure</code>) as the Direct MV Write Base Address for the Current Picture. <code>direct_mv_wr0_top</code> is selected for a frame picture or a top field picture, and <code>direct_mv_wr0_bottom</code> for a bottom field picture.</p> <p>The typical usage is to program this field to 16. And specifically, in order to support 16 reference frames, this field must be programmed as 16.</p> <p>For picture referencing/indexing, Frame Store ID Bit[6:0] = 7F (FF ???) is used to designate a non-existing picture; otherwise, Bit[7:5] should always be 0, For <code>img_dec_fs_idc[7:0]</code> (current decoding picture), it should never be a non-existing picture.</p> <p>(It was used to write the POC of the current frame to the POC list. Since we are now writing the entire POC list at a time this is no longer used for this purpose)</p> <p>In the current implementation, <code>img_dec_fs_idc[4:0]</code> can only be equal to 0 to 31. Bit[7:5] must be 0.</p> <p>For picture referencing/indexing, Frame Store ID Bit[6:0] = 7F (FF ???) is used to designate a non-existing picture; otherwise, Bit[7:5] should always be 0, For <code>img_dec_fs_idc[7:0]</code> (current decoding picture), it should never be a non-existing picture.</p> <p>(It was used to write the POC of the current frame to the POC list. Since we are now writing the entire POC list at a time this is no longer used)</p> <p>It is used to read the POC of the current frame for temporal direct motion vector weighting calculations. We could get the current POC with <code>POCList[RefList[0]]</code>, but instead use the existing <code>img_dec_fs_idc</code> to determine the POC from the <code>POCList[img_dec_fs_idc]</code>.</p> |
| 4 | 31:24 | <p>Residual Data Offset This gives the offset of the residual data in dwords. It should be 8 dword aligned. This should be programmed to the same value as the Residual Data Source Offset field in Vfe State Ex.</p> |
| | 23:18 | <p>Reserved: MBZ</p> |
| | 17 | <p>Thread Synchronization Overwrite. This field indicates whether hardware overwrites the Thread Synchronization field in the <code>MEDIA_OBJECT_EX</code> command of each output macroblock. When this field is set, the Thread Synchronization field is set only for intra macroblocks (and at least one of the neighbors A, B, C, D, Left-Bottom is available). Otherwise (if a macroblock is a P, B or I_PCM), the field is not set.</p> <p>0 = Not Overwrite 1 = Overwrite</p> |



| Dword | Bit | Description |
|-------|-------|--|
| | 16 | <p>Thread Synchronization. This field indicates whether the thread is a Synchronized Root Thread (SRT). When Thread Synchronization Overwrite is not set, this field is forwarded to the corresponding field in the MEDIA_OBJECT_EX command of each output macroblock. When Thread Synchronization Overwrite is set, hardware ignores this field.</p> <p>0 = Not a SRT 1 = SRT</p> |
| | 15:13 | Reserved. MBZ. |
| | 12 | <p>Enable16MV</p> <p>This field, when set, enables the 16MV motion vector format in the AVC-IT output data buffer. If this field is set to 0, all MvSize = 16MV are mapped to 32MV.</p> <p>0 – no MvSize of 16MV 1 – MvSize of 16MV enabled</p> |
| | 11:10 | <p>Chroma Format IDC, ChromaFormatIdc[1:0]</p> <p>It specifies the sampling of chroma component (Cb, Cr) in the current picture as follows :</p> <p>00 – monochrome picture 01 – 4:2:0 picture 10 – 4:2:2 picture (not supported) 11 – 4:4:4 picture (not supported)</p> <p>It is set to the value of the syntax element read from the current active SPS.</p> <p>The corresponding Monochrome Flag (monochrome_flag) can be derived from this field as follows :</p> <p>0 – Color Picture 1 – Monochrome picture (i.e. no Chroma components)</p> |
| | 9 | Reserved. MBZ. |
| | 8 | <p>Enable the In-Loop Deblocking Filter information write, EnableILDBWrite</p> <p>1 – enable ILDB writing output 0 – disable the ILDB writing output</p> <p>This bit controls ILDB information write to memory regardless of whether In-Loop Deblocking Filter is enabled or not for a given slice.</p> <p>If this bit is zero, even if a slice has in-loop deblocking enabled, no ILDB data is written to memory.</p> <p>On the other hand, if this bit is set, and if ILDB is disabled for a slice, ILDB data (default to no filtering) are written to memory for each macroblock in the slice. This allows ILDB to be performed at picture level regardless of the mixer of different slices (some has deblocking enabled and some disabled) within the picture. Default behavior when some slices of a picture has deblocking disabled. So, this will control the filling of those slices to the final output.</p> <p>This parameter is specified for Intel interface only.</p> |
| | 7 | <p>Entropy Coding Flag, entropy_coding_flag</p> <p>0 – CAVLD bit-serial decoding mode. 1 – CABAC bit-serial decoding mode.</p> <p>It specifies one of the two possible bit stream decoding modes in the AVC. It is set to the value of the syntax element read from the current active PPS.</p> |



| Dword | Bit | Description |
|-------|-----|--|
| | 6 | <p>Current Image Disposable Flag or Non-Reference Picture Flag, <code>ImgDisposableFlag</code></p> <p>0 – the current decoding picture may be used as a reference picture for others. 1 – the current decoding picture is not used as a reference picture (e.g. a B-picture cannot be a reference picture for any subsequent decoding)</p> <p>It is derived from <code>ImgDisposableFlag = (nal_ref_idc == 0)</code>. <code>nal_ref_idc</code> is a syntax element from a NAL unit.</p> |
| | 5 | <p>Constrained Intra Prediction Flag, <code>constrained_ipred_flag</code></p> <p>0 – allows both intra and inter neighbouring MB to be used in the intra-prediction decoding of the current MB. 1 – allows only to use neighboring Intra MBs in the intra-prediction decoding of the current MB. If the neighbor is an inter MB, it is considered as not available.</p> <p>It is set to the value of the syntax element in the current active PPS.</p> |
| | 4 | <p>Direct 8x8 Inference Flag, <code>direct_8x8_infer_flag</code></p> <p>0 – allows subpartitioning to go below 8x8 block size (i.e. 4x4, 8x4 or 4x8) 1 – allows processing only at 8x8 block size. MB Info is stored for 8x8 block size.</p> <p>It is set to the value of the syntax element in the current active SPS.</p> <p>It specifies the derivation process for luma motion vectors in the Direct MV coding modes (<code>B_Skip</code>, <code>B_Direct_16x16</code> and <code>B_Direct_8x8</code>). When <code>frame_mbs_only_flag</code> is equal to 0, <code>direct_8x8_inference_flag</code> shall be equal to 1.</p> <p>It must be consistent with the <code>frame_mbs_only_flag</code> and <code>transform_8x8_mode_flag</code> settings.</p> |
| | 3 | <p>8x8 IDCT Transform Mode Flag, <code>trans8x8_mode_flag</code></p> <p>Specifies 8x8 IDCT transform may be used in this picture</p> <p>0 – no 8x8 IDCT Transform, only 4x4 IDCT transform blocks are present 1 – 8x8 Transform is allowed</p> <p>It is set to the value of the syntax element in the current active PPS.</p> |
| | 2 | <p>Frame MB only flag, <code>frame_mbs_only_flag</code></p> <p>0 – not true ; effectively enables the possibility of MBAFF mode. 1 – true, only frame MBs can occur in this sequence, hence disallows the MBAFF mode and field picture.</p> <p>It is set to the value of the syntax element in the current active SPS.</p> |
| | 1 | <p>MBAFF mode is active, <code>mbaff_frame_flag</code>.</p> <p>This bit is valid only when the <code>img_structure[1:0]</code> indicates the current picture is a frame.</p> <p>0 – not in MBAFF mode 1 – in MBAFF mode</p> <p>It is derived from <code>MbaffFrameFlag = (mb_adaptive_frame_field_flag && !field_pic_flag)</code>. <code>mb_adaptive_frame_field_flag</code> is a syntax element in the current active SPS and <code>field_pic_flag</code> is a syntax element in the current Slice Header. They both are present only if <code>frame_mbs_only_flag</code> is 0. Although <code>mbaff_frame_flag</code> is a Slice Header parameter, its value is expected to be the same for all the slices of a picture.</p> <p>It must be consistent with the <code>mb_adaptive_frame_field_flag</code>, the <code>field_pic_flag</code> and the <code>frame_mbs_only_flag</code> settings.</p> |



| Dword | Bit | Description |
|-------|------|---|
| | 0 | Field picture flag, field_pic_flag , specifies the current slice is a coded field or not. 0 – a slice of a coded frame 1 – a slice of a coded field It is set to the same value as the syntax element in the Slice Header. It must be consistent (??? Expand out) with the img_structure[1:0] and the frame_mbs_only_flag settings. Although field_pic_flag is a Slice Header parameter, its value is expected to be the same for all the slices of a picture. |
| 5 | 31:0 | AVC-IT Command Header, AvcltCmdHeader It allows the driver to directly specify the DW0 of the AVC-IT Command to be generated by the BSD unit that will pass down to drive the VFE unit. This parameter is specified for Intel interface only. |

2.7.2.2 AVC_BSD_QM_STATE Command

This command is only issued when qm_present_flag is set in the IMG_STATE command. If qm_present_flag is not set in the IMG_State command, and AVC_BSD_QM_STATE command is issued, the BCS will still process the command, but the BSD unit will ignore the data being forwarded.

Quantization matrices are loaded only if

- 1) the current Slice is the first of a picture,
- 2) active SPS or active PPS has read in,
- 3) active PPS has changed between pictures

Only the latest scaling matrices read from the bitstream in processing the current active SPS and PPS are sent down to the BSD unit, as needed. Maximum there can be 8 matrices. If transform_8x8_mode_flag is set, maximum two 8x8 scaling matrices and six 4x4 scaling matrices may be present; otherwise only maximum six 4x4 scaling matrices may be present. Hence, this state descriptor has a variable length; the Present Flag, the List_Present flags and the Use_Default flags together can determine which and how many scaling matrices to follow. If the qm_present_flag is set in the IMG_State command, but a use_default flag of QM_State command is set, the corresponding scaling matrix is not sent. That is, all default scaling matrices are generated inside the BSD unit.

Each entry of a scaling matrix occupies 1 byte. Hence, the minimum size of the state descriptor is 1 Dword, and the maximum is 57 Dwords.

The Scaling Matrices, if present, must follow a strict ordering in the state descriptor. They are packed in the ascending order of the i index. Hence, the Dword 1 is referring to the Scaling Matrix with the lowest i. This state descriptor may be sent prior to decoding each Slice. The values set for these state variables are retained internally across Slices (QM_State is actually separated out from the IMG_State), until they are reset by H/W Asynchronous Reset or changed by the next AVC_BSD_QM_STATE Command.



| Dword | Bit | Description |
|---------------|----------|---|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | AVC Command Opcode = AVC_BSD_QM_STATE Pipeline[28:27] = BSD = 2h; Opcode[26:24] = AVC = 4h; Sub Opcode[23:16] = 1h |
| | 15:0 | DWord Length (Excludes DWords 0,1) = 0000h to max 0038h |
| 1 | 31:16 | Reserved. MBZ |
| | 15:8 | Use built-in Default QM Flags for the current Slice, use_default_flags[i], i = 0 to 7 i=0 – 4x4 Intra Y i=1 – 4x4 Intra Cr i=2 – 4x4 Intra Cb i=3 – 4x4 Inter Y i=4 – 4x4 Inter Cr i=5 – 4x4 Inter Cb i=6 – 8x8 Intra Y i=7 – 8x8 Inter Y use_default_flags[i] = 1, and if the corresponding qm_list_flags[i] and qm_present_flag are also set, then use the built-in default scaling matrix use_default_flags[i] = 0, indicates the built-in default QM matrix is not used, and the scaling matrix read from the bitstream is used. qm_present_flag (in IMG_STATE), qm_list_flags[i], use_default_flags[i] and LevelScale matrices must be set in consistent and coherent. This parameter is specified for Intel interface only. |
| | 7:0 | QM List Present Flags for the current Slice, qm_list_flags[i], i = 0 to 7 i=0 – 4x4 Intra Y i=1 – 4x4 Intra Cb i=2 – 4x4 Intra Cr i=3 – 4x4 Inter Y i=4 – 4x4 Inter Cb i=5 – 4x4 Inter Cr i=6 – 8x8 Intra Y i=7 – 8x8 Inter Y qm_list_flags[i] = 1 and qm_present_flag is set, indicates either using the scaling matrix read from the bit stream, or use the built-in default matrix if the Use_Default is also set. The following LevelScale field is present if and only if the corresponding bit in this field is set. qm_list_flag[i] = 0, use the fallback rule specified in the AVC spec. qm_present_flag (in IMG_STATE), qm_list_flags[i], use_default_flags[i] and LevelScale matrices must be set in consistent and coherent. This parameter is specified for Intel interface only. |
| 0 or 4 Dwords | 16 bytes | Luma4x4 Intra WeightScale (i=0) Unsigned integer value, ranging from 0 to 255. An array of 4x4 scaling values (one-to-one correspondence with the coefficient position) for an Intra Luma block that is stored in raster order. It is set to the values derived from the syntax elements in the current active PPS and active SPS. |



| Dword | Bit | Description |
|----------------|----------|--|
| 0 or 4 Dwords | 16 bytes | <p>Cb4x4 Intra WeightScale (i=1) Unsigned integer value, ranging from 0 to 255. An array of 4x4 scaling values (one-to-one correspondence with the coefficient position) for an Intra Chroma Cb block that is stored in row major order (i.e. raster scan order). It is set to the values derived from the syntax elements in the current active PPS and active SPS.</p> |
| 0 or 4 Dwords | 16 bytes | <p>Cr4x4 Intra WeightScale (i=2) Unsigned integer value, ranging from 0 to 255. An array of 4x4 scaling values (one-to-one correspondence with the coefficient position) for an Intra Chroma Cr block that is stored in row major order (i.e. raster scan order). It is set to the values derived from the syntax elements in the current active PPS and active SPS.</p> |
| 0 or 4 Dwords | 16 bytes | <p>Luma4x4 Inter WeightScale (i=3) Unsigned integer value, ranging from 0 to 255. An array of 4x4 scaling values (one-to-one correspondence with the coefficient position) for an Inter Luma block that is stored in row major order (i.e. raster scan order). It is set to the values derived from the syntax elements in the current active PPS and active SPS.</p> |
| 0 or 4 Dwords | 16 bytes | <p>Cb4x4 Inter WeightScale (i=4) Unsigned integer value, ranging from 0 to 255. An array of 4x4 scaling values (one-to-one correspondence with the coefficient position) for an Inter Chroma Cb block that is stored in row major order (i.e. raster scan order). It is set to the values derived from the syntax elements in the current active PPS and active SPS.</p> |
| 0 or 4 Dwords | 16 bytes | <p>Cr4x4 Inter WeightScale (i=5) Unsigned integer value, ranging from 0 to 255. An array of 4x4 scaling values (one-to-one correspondence with the coefficient position) for an Inter Chroma Cr block that is stored in row major order (i.e. raster scan order). It is set to the values derived from the syntax elements in the current active PPS and active SPS.</p> |
| 0 or 16 Dwords | 64 bytes | <p>Luma8x8 Intra WeightScale (i=6) Unsigned integer value, ranging from 0 to 255. An array of 8x8 scaling values (one-to-one correspondence with the coefficient position) for an Intra Luma block that is stored in row major order (i.e. raster scan order). It is set to the values derived from the syntax elements in the current active PPS and active SPS.</p> |
| 0 or 16 Dwords | 64 bytes | <p>Luma8x8 Inter WeightScale (i=7) Unsigned integer value, ranging from 0 to 255. An array of 8x8 scaling values (one-to-one correspondence with the coefficient position) for an Inter Luma block that is stored in row major order (i.e. raster scan order). It is set to the values derived from the syntax elements in the current active PPS and active SPS.</p> |



2.7.2.3 AVC_BSD_SLICE_STATE Command

This state descriptor includes all the information for

1. Inter-Prediction Reference Lists L0 and L1
2. Weights and Offset for Inter-Prediction

They are all syntax elements read from or derived from the Slice Header and may, therefore, be changed on a Slice boundary.

The state descriptor of the Slice_State command is of variable size. However, all the constitute components, when present, are of fixed size structure. The Slice_State command should not be issued at all, if there is no inter-prediction decoding or none of its components are present. If the present flag is set to 0, and the corresponding data is still sent, unpredicted result. If

The state information for motion vector prediction specifies the frame indices of the reference list list_0 and list_1. It is sent to the MPR unit for determining which frames to fetch for Temporal Direct Motion Vectors. These MPR states may be changed on a Slice boundary, the entire reference lists are sent, and are packed in a consecutive order (and there is no bubble – i.e. no empty/invalid/unspecified entry in between).

The very first entry of each reference list is always started with index 0, and is named as the top of the reference list.

For the weights and offsets specification in the SLICE_STATE state descriptor, it specifies the weights and offsets for inter-prediction using the reference list list_0 and list_1. It is outputted from the BSD unit to the AVC-IT interface. These Weights and Offsets states are set to the values of syntax elements read from the Slice Header, which may be changed on a Slice boundary. They are packed in a consecutive order (and there is no bubble – i.e. no empty/invalid/unspecified entry in between) . As a result, the weight and offset state descriptor has a fixed length, if it is present.

Only when WeightedPredFlag = 1 (enable weighted prediction), will there be a L0 weight+ offset table (and there is no L1 stuff in P slice) being sent to the BSD unit through the Slice_State command. If WeightedBiPredIdc[1:0] != 1 and WeightedPredFlag = 0, no Slice_State command should be issued for this table.

Only when WeightedBiPredIdc[1:0] = 1 (explicit weighted prediction), will there be a L1 and/or a L0 weight+offset tables being sent to the BSD unit through the Slice_State command. If WeightedBiPredIdc[1:0] != 1 and WeightedPredFlag = 0, no Slice_State command should be issued for this table.

As a result, this state descriptor has a variable length, and is predetermined by the driver. The minimum size is 2 Dwords if this command is ever issued, since at minimum even there is no L1 and L0 and no weight and offset, there is still DW0 and DW1. The maximum size is 113 Dwords.

Weight-Offset programming model[Errata – Dev CTG] Driver needs to look ahead all slices for the frame and check if any weight value is 128. This can only happen when luma_weight_10/11/chroma_weight_10/11 flags are 0. If it finds a 128 then it needs to remap this value to a non used value within a 16 bit twos compliment range and then give the remapped number to the kernel for identifying this case.



| Dword | Bit | Description |
|-------|---|--|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | AVC Command Opcode = AVC_BSD_SLICE_STATE Pipeline[28:27] = BSD = 2h; Opcode[26:24] = AVC = 4h; Sub Opcode[23:16] = 2h |
| | 15:0 | DWord Length (Excludes DWords 0,1) = 0000h (min) to 00D0h (max) |
| 1 | 31:4 | Reserved. MBZ. |
| | 3 | Weight+Offset L1Table Present Flag If Weight+OffsetL1 Table Present Flag = 1, indicates the full 32-entry L1 Table for Y, Cb and Cr is followed; otherwise it is not present at all (no further inline data). It must be set in consistent with the WeightedPredFlag and WeightedBiPredIdc in the Img_State command. This parameter is specified for Intel interface only. |
| | 2 | Weight+OffsetL0 Table Present Flag If Weight+OffsetL0 Table Present Flag = 1, indicates the full 32-entry L0 Table for Y, Cb and Cr is followed; otherwise it is not present at all (no further inline data). It must be set in consistent with the WeightedPredFlag and WeightedBiPredIdc in the Img_State command. This parameter is specified for Intel interface only. |
| | 1 | RefPicListL1 Present Flag If RefPicListL1 Present Flag = 1, indicates the full 32-entry L1 list is followed; otherwise it is not present at all (no further inline data). It must be set in consistent with the Num_Ref_Idx_L1 derived from the syntax element in the Slice Header, and if it is a B Slice. This parameter is specified for Intel interface only. |
| 0 | RefPicListL0 Present Flag If RefPicListL0 Present Flag = 1, indicates the full 32-entry L0 list is followed; otherwise, it is not present at all (no further inline data). It must be set in consistent with the Num_Ref_Idx_I0 derived from the syntax element in the Slice Header, and if it is a P or B Slice. This parameter is specified for Intel interface only. | |



| Dword | Bit | Description |
|----------------|--------------|--|
| 2-9 or none | 31:0 each | <p>Reference List0 Write</p> <p>This field is present only if the RefPicListL0 Present Flag = 1; otherwise, it should not be present.</p> <p>If present, it always specifies 32 reference pictures in list0, regardless they are “existing picture” or not. If a picture is non-existing, the corresponding entry should be set to all ones.</p> <p>Each list entry is 1 byte. A 32-bit DW can hold 4 list entries in the following format</p> <p style="padding-left: 20px;">31:24 entry X+3 (e.g. list0_3)</p> <p style="padding-left: 20px;">23:16 entry X+2 (e.g. list0_2)</p> <p style="padding-left: 20px;">15:8 entry X+1 (e.g. list0_1)</p> <p style="padding-left: 20px;">7:0 entry X (e.g. list0_0)</p> <p style="padding-left: 20px;">X is replaced by the paddr[2:0] * 4 ; paddr[5:0] with 0x30 and 0x37</p> <p>The byte definition for a reference picture :</p> <p>Bit 7 : Non-Existing – indicates that frame store index that should have been at this entry did not exist and was replaced by an index for error concealment</p> <p>Bit 6 : Long term bit – set this reference picture to be used as long term reference</p> <p>Bit 5 : Field picture flag – indicates frame/field</p> <p>Bit 4:0 : Frame store index or Frame Store ID (Bit 4:1 is used to form the binding table index in Intel implementation)</p> <p>This is the final Reference List L0 after any reordering specified in the Slice Header as well as modified by the driver, and its indices values are all translated to the Intel specification.</p> <p>If the reference picture is a frame (Bit5 = 1), frame store ID is always an even number.</p> <p>The frame store ID = 0 and 1 are assigned to the current decoding picture; hence reference picture ID will start from 2 onward.</p> <p>This list is used in MV information output of the BSD unit (but non-existing picture are mapped to 00 instead to smilify the kernel). DMV also read and write Mvlist0 using this frame store ID.</p> |



| Dword | Bit | Description |
|------------------|--------------|--|
| 10-17 or none | 31:0 each | <p>Reference List1 Write</p> <p>This field is present only if the RefPicListL1 Present Flag = 1; otherwise, it should not be present.</p> <p>If present, it always specifies 32 reference pictures in list1, regardless they are existing or not. If a picture is non-existing, the corresponding entry should be set to all ones.</p> <p>Each list entry is 1 byte. A 32-bit DW can hold 4 list entries in the following format</p> <p style="padding-left: 40px;">31:24 entry X+3 (e.g. list1_7)</p> <p style="padding-left: 40px;">23:16 entry X+2 (e.g. list1_6)</p> <p style="padding-left: 40px;">15:8 entry X+1 (e.g. list1_5)</p> <p style="padding-left: 40px;">7:0 entry X (e.g. list1_4)</p> <p style="padding-left: 40px;">X is replaced by the paddr[2:0] * 4 ; paddr[5:0] with 0x38 and 0x3F</p> <p>The byte definition for a reference picture:</p> <p>Bit 7 : Non-Existing – indicates that frame store index that should have been at this entry did not exist and was replaced by an index to for error concealment</p> <p>Bit 6 : Long term bit – set this reference picture to be used as long term reference</p> <p>Bit 5 : Field picture flag – indicates frame/field</p> <p>Bit 4:0 : Frame store index (or also known as binding table index in Intel implementation)</p> <p>This is the final Reference List L1 after any reordering specified in the Slice Header as well as modified by the driver, and its indices values are all translated to the Intel specification.</p> <p>This list is used in MV information output of the BSD unit. DMV also read and write Mvlist1 using this frame store ID.</p> |



| Dword | Bit | Description |
|------------------|--------------|--|
| 18-65 or none | 31:0 each | <p>WeightOffset[L0=0][i=0 to 31][Y/Cb/Cr][weight/offset]</p> <p>WeightOffset[0][i=0][Y=0][Offset=0]</p> <p>WeightOffset[0][i=0][Y=0][Weight=1]</p> <p>WeightOffset[0][i=0][Cb=1][Offset=0]</p> <p>WeightOffset[0][i=0][Cb=1][Weight=1]</p> <p>WeightOffset[0][i=0][Cr=2][Offset=0]</p> <p>WeightOffset[0][i=0][Cr=2][Weight=1]</p> <p>:</p> <p>WeightOffset[0][i=31][Y=0][Offset=0]</p> <p>WeightOffset[0][i=31][Y=0][Weight=1]</p> <p>WeightOffset[0][i=31][Cb=1][Offset=0]</p> <p>WeightOffset[0][i=31][Cb=1][Weight=1]</p> <p>WeightOffset[0][i=31][Cr=2][Offset=0]</p> <p>WeightOffset[0][i=31][Cr=2][Weight=1]</p> <p>Signed integer values, ranging from -128 to 127.</p> <p>This field is present only if the Weight+OffsetL0 Table Present Flag= 1; otherwise, it should not be present. If present, the full table is always specified. Any reference list L0[i] that does not exist, the corresponding weight and offset are set to 0.</p> <p>Weight and Offset are 1 byte each. Weight is the high byte and Offset is the lower byte. Hence, we are packing 2 sets of weight and offset into 1 DW. Bits[31:16] will store the set of weight and offset will higher values of i or Y/Cr/Cb (represented as 0/1/2 in the indexing).</p> <p>WeightOffset[L0=0][i=0 to 31][Y=0] (i.e. luma_weight_I0[i]) are specified for the weighting and offset factors applied to the luma prediction value for list 0 prediction using RefPicList0[i] (one-to-one correspondence in i). When luma_weight_I0_flag (Slice Header syntax element) is equal to 1, the value of luma_weight_I0[i] shall be in the range of -128 to 127. When luma_weight_I0_flag is equal to 0, luma_weight_I0[i] shall be inferred to be equal to $2^{\text{luma_log2_weight_denom}}$ for RefPicList0[i]. luma_log2_weight_denom is a Slice Header syntax element.</p> <p>WeightOffset[L0=0][i=0 to 31][Cb=1] (i.e. chromaCb_weight_I0[i]) are specified for the weighting and offset factors applied to the chroma Cb prediction values for list 0 prediction using RefPicList0[i] (one-to-one correspondence in i). When chroma_weight_I0_flag (Slice Header syntax element) is equal to 1, the value of chromaCb_weight_I0[i] shall be in the range of -128 to 127. When chroma_weight_I0_flag is equal to 0, chromaCb_weight_I0[i] shall be inferred to be equal to $2^{\text{chroma_log2_weight_denom}}$ for RefPicList0[i]. chroma_log2_weight_denom is a Slice Header syntax element.</p> <p>WeightOffset[L0=0][i=0 to 31][Cr=2] (i.e. chromaCr_weight_I0[i]) are specified for the weighting and offset factors applied to the chroma Cr prediction values for list 0 prediction using RefPicList0[i] (one-to-one correspondence in i). When chroma_weight_I0_flag (Slice Header syntax element) is equal to 1, the value of chromaCr_weight_I0[i] shall be in the range of -128 to 127. When chroma_weight_I0_flag is equal to 0, chromaCr_weight_I0[i] shall be inferred to be equal to $2^{\text{chroma_log2_weight_denom}}$ for RefPicList0[i].</p> <p>The table is set to the values derived from the syntax elements read in the Slice Header.</p> |



| Dword | Bit | Description |
|-------------------|--------------|--|
| 66-113 or none | 31:0 each | <p>WeightOffset[L1=1][32][Y/Cb/Cr][weight/offset]</p> <p>WeightOffset[1][0][Y=0][Offset=0]</p> <p>WeightOffset[1][0][Y=0][Weight=1]</p> <p>WeightOffset[1][0][Cb=1][Offset=0]</p> <p>WeightOffset[1][0][Cb=1][Weight=1]</p> <p>WeightOffset[1][0][Cr=2][Offset=0]</p> <p>WeightOffset[1][0][Cr=2][Weight=1]</p> <p>:</p> <p>WeightOffset[1][31][Y=0][Offset=0]</p> <p>WeightOffset[1][31][Y=0][Weight=1]</p> <p>WeightOffset[1][31][Cb=1][Offset=0]</p> <p>WeightOffset[1][31][Cb=1][Weight=1]</p> <p>WeightOffset[1][31][Cr=2][Offset=0]</p> <p>WeightOffset[1][31][Cr=2][Weight=1]</p> <p>Signed integer values, ranging from -128 to 127.</p> <p>This field is present only if the Weight+OffsetL1 Table Present Flag= 1; otherwise, it should not be present. If present, the full table is always specified. Any reference list L1[i] that does not exist, the corresponding weight and offset are set to 0.</p> <p>Weight and Offset are 1 byte each. Weight is the high byte and Offset is the lower byte. Hence, we are packing 2 sets of weight and offset into 1 DW. Bits[31:16] will store the set of weight and offset will higher values of i or Y/Cr/Cb (represented as 0/1/2 in the indexing).</p> <p>WeightOffset[L1=1][i=0 to 31][Y=0] (i.e. luma_weight_11[i]) are specified for the weighting and offset factors applied to the luma prediction value for list 1 prediction using RefPicList1[i] (one-to-one correspondence in i). When luma_weight_11_flag (Slice Header syntax element) is equal to 1, the value of luma_weight_11[i] shall be in the range of -128 to 127. When luma_weight_11_flag is equal to 0, luma_weight_11[i] shall be inferred to be equal to $2^{\text{luma_log2_weight_denom}}$ for RefPicList1[i]. luma_log2_weight_denom is a Slice Header syntax element.</p> <p>WeightOffset[L1=0][i=0 to 31][Cb=1] (i.e. chromaCb_weight_11[i]) are specified for the weighting and offset factors applied to the chroma Cb prediction values for list 1 prediction using RefPicList1[i] (one-to-one correspondence in i). When chroma_weight_11_flag (Slice Header syntax element) is equal to 1, the value of chromaCb_weight_11[i] shall be in the range of -128 to 127. When chroma_weight_11_flag is equal to 0, chromaCb_weight_11[i] shall be inferred to be equal to $2^{\text{chroma_log2_weight_denom}}$ for RefPicList1[i]. chroma_log2_weight_denom is a Slice Header syntax element.</p> <p>WeightOffset[L1=0][i=0 to 31][Cr=2] (i.e. chromaCr_weight_11[i]) are specified for the weighting and offset factors applied to the chroma Cr prediction values for list 1 prediction using RefPicList1[i] (one-to-one correspondence in i). When chroma_weight_11_flag (Slice Header syntax element) is equal to 1, the value of chromaCr_weight_11[i] shall be in the range of -128 to 127. When chroma_weight_11_flag is equal to 0, chromaCr_weight_11[i] shall be inferred to be equal to $2^{\text{chroma_log2_weight_denom}}$ for RefPicList1[i].</p> <p>The table is set to the values derived from the syntax elements read in the Slice Header.</p> |



2.7.2.4 AVC_BSD_BUF_BASE_STATE Command

This command contains the base addresses for memory buffers allocated for BSD operations. It does not include the byte stream read address, which is to be specified in the AVC_BSD_OBJECT Command.

There are six base addresses being defined:

1. BSD Row Store memory base address for R/W,
2. MPR Row Store memory base address for R/W,
3. MB Info memory base address for Write-only,
4. Decoded Uncompressed Coefficients, IPCM Coefficients or MV array memory base address for Write-only,
5. Direct MV memory base address for Write-only, and
6. 31 Direct MV memory base addresses for Read-only.
 - a. Not including the current frame, which is represented by the Write Base address
 - b. Although the number of Direct MV base addresses should be the same as the number of reference frames in the reference list (specified in the Slice_STATE), since it is possible to have the same frame presents in both list0 and list1 and there is no logics to cross checking, it is decided to send the complete addressess list always.

The state descriptor provides the state and information required to decode a Slice, and their values may subject to change on a Slice boundary. It should be sent prior to decoding each Slice. The values set for these state variables are retained internally across Slices, until they are reset by H/W Asynchronous Reset or changed by the next AVC_BSD_BUF_BASE_STATE Command.

When switching video stream, these base address must be reprogrammed.

Xxx

| Dword | Bit | Description |
|-------|-------|---|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | AVC Command Opcode = AVC_BSD_BUF_BASE_STATE Pipeline[28:27] = BSD = 2h; Opcode[26:24] = AVC = 4h; Sub Opcode[23:16] = 3h |
| | 15:0 | DWord Length (Excludes DWords 0,1) = 0048h |
| 1 | 31:6 | BSD Row Store Base Address This field provides the base address of the scratch buffer used by BSD unit to store BSD information for the previous row to assist BSD decoding of the macroblocks in the current row. The BSD Row Store buffer must be 64-byte cacheline aligned. 1.5 cacheline (CL) per MB when in MBAFF mode (row of MB pair); 0.75 CL per MB for non-MBAFF. So, to support 120 MBs per row (1920 screen resolution), 1.5 * 120 * 64 bytes = 11,520 bytes are required. Half cacheline alignment should be followed. Hardware uses the horizontal address of each macroblock to address the BSD Row Store. |
| | 5:0 | Reserved : MBZ |



| Dword | Bit | Description |
|-------|-------|---|
| 2 | 31:6 | <p>MPR Row Store Base Address</p> <p>This field provides the base address of the scratch buffer used by BSD unit to store MPR information for the previous row to assist MPR decoding of the macroblocks in the current row.</p> <p>1 cacheline (CL) per MB when in MBAFF mode (row of MB pair); 0.5 CL per MB for non-MBAFF, So, ot support 120 MBs per row (1920 screen resolution), $1 * 120 * 64$ bytes = 7,680 bytes are required. Half cacheline alignment should be followed.</p> <p>Hardware uses the horizontal address of each macroblock to address the MPR Row Store.</p> |
| | 5:0 | Reserved : MBZ |
| 3 | 31:6 | <p>AVC-IT Command MB Info Write Base Address (defined for streamout mode)</p> <p>This field provides the base address of the AVC-IT command output buffer that contains the MB Info needed for inverse transform and motion compensation. This buffer must be 64-byte cacheline aligned.</p> <p>Hardware uses the MB number within the picture to address this buffer..</p> <p>[Also provide the data structure for each macroblock. This is important as the interface is shared between this pipe and the 3DPIPE.]</p> <p>[Addressing in raster order and also the wave-front order. Providing a reference to the description of these two orders. Also discuss how to handle MBAFF.]</p> <p>1 cacheline (CL) for each MB, 8160 MBs (1920x1088 screen resolution) per frame, hence $1 * 8160 * 64$ bytes per frame = 522,240 bytes are required. Cacheline alignment should be followed.</p> |
| | 5:0 | Reserved : MBZ |
| 4 | 31:12 | <p>AVC-IT Data Write Base Address</p> <p>This is the 4KB aligned portion of the base address for the AVC-IT data write buffer. Together with the AVC-IT Data Write Offset, it forms the base pointer where the AVC-IT data (including for example motion vectors, weight-offsets, non-zero residual data), where data for the first macroblock will be written to.</p> <p>This field affect the final memory address computed for the AVC-IT Data output. However, it doesn't affect the Indirect Data Start Address computed by BSD hardware that is in the AVC-IT Command output (DW3 of the output AVC-IT Command).</p> <p>Each macroblock within the AVC-IT data buffer has a fixed size even though only valid data are written to memory. The remaining bytes for each macroblock are undefined. Assuming 32 cachelines are needed per macroblock, a buffer (with a resolution of 1920x1088) is 16,711,680 bytes.</p> <p>Format = GraphicsAddress[31:12]</p> <p><i>Programming Note: As the AVC-IT data buffer will be used by GPE as the indirect object buffer for MEDIA_OBJECT_EX command, this field should be programmed the same as the Indirect Object Base Address field of the STATE_BASE_ADDRESS command.</i></p> |
| | 11:6 | <p>AVC-IT Data Write Offset</p> <p>This offset is relative to the AVC-IT Data Write Buffer Base Address. This field, together with the AVC-IT Data Write Base Address, forms the 64-byte cacheline aligned address of the AVC-IT data write buffer.</p> <p>This field is used to compute the Indirect Data Start Address in the AVC-IT Command output (DW3 of the output AVC-IT Command).</p> |



| Dword | Bit | Description |
|-------|------|--|
| | 5:0 | Reserved : MBZ |
| 5 | 31:6 | <p>ILDB Data Write Base Address</p> <p>This field provides the base address of the ILDB output data buffer. This buffer must be 64-byte cacheline aligned.</p> <p>2 cachelines (CL) for each MB, 8160 MBs (1920x1088 screen resolution) per frame, hence 2*8160*64 bytes per frame = 1,044,480 bytes are required. Real data only in 1 CL, extra CL added due to driver/kernel requirement for data expansion. Cacheline alignment should be followed.</p> |
| | 5:0 | Reserved : MBZ |
| 6 | 31:6 | <p>Direct MV Read Base Address for Reference Frame 0, 64-bytes cache-line aligned – direct_mv_rd0[31:6]</p> <p>This is a ROW Store private buffer space.</p> <p>The read buffer size is 557,056 bytes for 1 frame (the selected colPic). It is scalable with frame height, but does not scale with frame width as H/W uses a fixed with of 128 MBs for row store buffers which is the smallest power of 2 value larger than 120 – 1920x1088 screen resolution.</p> <p>This field may also be used as the Direct MV Write Base Address for the Current Picture if img_dec_fs_idc is set to 0 to 15. See img_dec_fs_idc for more details.</p> |
| | 5:0 | Reserved : MBZ |
| | 5:0 | Reserved : MBZ |
| 7-37 | 31:6 | <p>Direct MV Read Base Address for Reference Frame 1 to Reference Frame 31</p> <p>Definition follows that of direct_mv_rd0[31:6]</p> |
| | 5:0 | Reserved : MBZ |
| 38 | 31:6 | <p>Direct MV Write Base Address for the Current Picture (default to be Reference Frame 0 TOP field), 64-bytes cache-line aligned – direct_mv_wr0_top[31:6]</p> <p>This is a ROW Store private buffer space</p> <p>The write buffer size is 557,056 bytes for 1 frame (the selected colPic). It is scalable with frame height, but does not scale with frame width as H/W uses a fixed with of 128 MBs for row store buffers which is the smallest power of 2 value larger than 120 – 1920x1088 screen resolution.</p> <p>There is only one write buffer for capturing the DMVs for the current picture.</p> <p>This field is used if img_dec_fs_idc is 16 and the current picture is a frame picture or a top field picture.</p> |
| | 5:0 | Reserved : MBZ |



| Dword | Bit | Description |
|-------|------|---|
| 39 | 31:6 | <p>Direct MV Write Base Address for the Current Picture (default to be Reference Frame 0 BOTTOM field), 64-bytes cache-line aligned – direct_mv_wr0_bottom[31:6]</p> <p>This is a ROW Store private buffer space</p> <p>The write buffer size is 557,056 bytes for 1 frame (the selected colPic). It is scalable with frame height, but does not scale with frame width as H/W uses a fixed width of 128 MBs for row store buffers which is the smallest power of 2 value larger than 120 – 1920x1088 screen resolution.</p> <p>There is only one write buffer for capturing the DMVs for the current picture.</p> <p>This field is used if img_dec_fs_idc is 16 and the current picture is a bottom field picture.</p> |
| | 5:0 | Reserved : MBZ |
| 40-73 | 31:0 | <p>POC List, POCList[0...33][31:0]</p> <p>Each POC value is a signed 32-bit number.</p> <p>One-to-one correspondence with the 34 Direct MV Read/Write Address for Reference Frames.</p> <p>There are 34 POC entries in the list, indexed by the frame_store_ID. POCList[0] is the CurrPic POC. If it is in a field mode, POCList[0] is the CurrFOC for the top field and the POCList[1] is the CurrFOC for the bottom field.</p> <p>For frame mode, every other entry is skipped.</p> <p>Frame_Store_ID[5:1]+T/B bit[0] is used to index into the POCList.</p> |

2.7.3 AVC_BSD_OBJECT Command

The AVC_BSD_OBJECT command is the only primitive command for the AVC_BSD Unit. The same command is used for both CABAC and CAVLD modes. The Slice Data portion of the bitstream is loaded as indirect data object.

Before issuing an AVC_BSD_OBJECT command, all BSD states need to be valid. Therefore the commands used to set these states need to have been issued at some point prior to the issue of an AVC_BSD_OBJECT command.

The Encryption of bit stream is indicated in bit 31 of DW 1. For such encrypted bit streams, this command has one extra dword (DW 8), to load in the initial counter value for AES-128 Counter mode decryption of the bit stream.

[DevILK] supports transport packet based encryption. Bit 30 of DW1 indicates such packet based encryption mode. Bit 29 of DW 1 specifies the packet size. For packet-based encrypted bitstream, this command has valid information in three additional dwords (DW 9 through 11).

To handle encrypted bitstream decoding, host software is required to align the 16byte chunk containing this first byte of actual bit-stream slice, to a naturally aligned 16byte boundary, i.e. that 16byte chunk should start on an aligned 16byte address boundary. Host software should ensure that after the last valid byte of bitstream data and before encryption, enough padding is added up to the end of a 16byte chunk. In addition, for a protected context, all kernel source and destination data surface addresses must be in PCM space.



| Dword | Bits | Description |
|-------|-------|--|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | AVC Command Opcode = AVC_BSD_OBJECT Pipeline[28:27] = BSD = 2h; Opcode[26:24] = AVC = 4h; Sub Opcode[23:16] = 8h |
| | 15:0 | DWord Length (Excludes DWords 0,1) [DevCTG] = 0006h without encryption (bit 31 of DW1 is set to 0), where inline data are in dwords 3-7 = 0007h with AES decryption on (bit 31 of DW1 is set to 1), where inline data are in dwords 3-8 [DevILK] = 000Ah for all cases where inline data are in dwords 3-15 |
| 1 | 31 | Bit stream is AES-128 Counter mode Encrypted '0' – Not encrypted '1' - Encrypted |
| | 30 | [DevCTG] Reserved – MBZ [DevILK] Packet-Based Bit Stream. The input bitstream data is in transport packet format with potential gaps within each 192-Byte packet. '0' – Not packet-based (bitstream data is a contiguous sequence of bits) '1' - Packet-based This bit is a don't care if the bit stream itself is not encrypted, i.e. bit_31 = '0'. |
| | 29 | [DevCTG, Reserved – MBZ] [DevILK] Packet Format. Valid information, only for a packet-based bitstream (bit[30]) is set to 1. Else, this bit is a don't care. '0' – Reserved (was 2KB packet format). '1' – 192B packet format. This information is used to determine the packet size of the bitstream generated. Header information is at the beginning of each packet. |
| | 28:22 | Reserved – MBZ |
| | 21:0 | Indirect Data Length. This field provides the length in bytes of the indirect data. A value zero indicates that indirect data fetching is disabled – subsequently, the Indirect Data Start Address field is ignored. This field must have the same alignment as the Indirect Object Data Start Address. It is the length in bytes of the bitstream data for the current slice. It includes the first byte of the first macroblock and the last byte of the last macroblock in the slice. Specifically, the zero-padding bytes (if present) and the next start-code are excluded. Hardware ignores the contents after the last non-zero byte. This field is sized to support AVC High Profile 4.1 Level bitstream. Interpreted from the AVC Spec, the maximum number of bits per macroblock for 4:2:0 is 3072 (bounded by IPCM (16x16+64x2)x8). So the maximum slice size (e.g. for 1080i) is 3072 x120 x 68 / 8 = 3133440 bytes, which requires 22 bits. It includes the byte that contains the First MB Bit Offset In implementing a phantom slice at the end of a picture for automatic error concealment, this field should set to 0. For packet-based decryption, this byte length does not include the length for header bytes, but only the byte length for actually valid bitstream slice information. Format = U22 in bytes |



| Dword | Bits | Description |
|------------|--------------|---|
| 2 | 31:29 | Reserved : MBZ |
| | 28:0 | <p>Indirect Data Start Address. This field specifies the Graphics Memory starting address of the data to be fetched into BSD Unit for processing. This pointer is relative to the AVC Indirect Object Base Address.</p> <p>Hardware ignores this field if indirect data is not present.</p> <p>It is a byte-aligned address for the AVC bitstream data in both CABAC/CAVLD Modes.</p> <p>Note however for encrypted bit stream, the 16byte chunk containing this first byte of actual bitstream slice, should be naturally aligned, i.e. that 16byte chunk should start on an aligned 16byte address boundary. In other words the first byte of that 16byte chunk should be aligned to a 16byte boundary.</p> <p>In implementing a phantom slice at the end of a picture for automatic error concealment, this field should set to 0. Range = [0 - 512MB]</p> |
| 3-variable | 31:0 Each | <p>Inline Data</p> <p>All the required Slice Header parameters and error handling settings are captured as inline data of the AVC_BSD_OBJECT command. The length is device dependent. Its definition is described in the next section.</p> |

Note that scattered raw slice data format is not supported by hardware. The size of a compression buffer for a slice needs to be allocated, which can be big, according to the minimal compression ratio required by a given profile/level. Hardware doesn't support the option for the driver to allocate 'nominally sized raw slice buffers' and use multiple such buffers for the worst case slice data.

2.7.3.1 Inline Data Description

The Inline Data includes all the required slice decoding states, extracted primarily from the Slice Header and its derivatives. It provides information for the following operations:

1. CABAC/CAVLD decoding
2. Internal error handling at the Slice boundary
 - a. H/W Automatic Error Concealment
3. Motion vector prediction decoding (MPR)
4. BSD output compositing (feeding the subsequent IT/MC/ILDB operations)

The only H/W error detection and concealment mechanism is comparing the Slice_Start_MB_Num of decoding the new current slice with the Current_MB_Num resulted from decoding the previous slice. If they do not match, an error is assumed. There are two possible cases, either "greater than" or "less than". If the Slice_Start_MB_Num is less than the Current_MB_Num, the Current_MB_Num is adjusted internally to be the same as the Slice_Start_MB_Num. If the Slice_Start_MB_Num is greater than, there is a gap between the Current_MB_Num and the Slice_Start_MB_Num that needs to be filled. The filling is started from [Current_MB_Num – Rewind_Num], and the method of filling can be either intra or inter as specified in the state descriptor.

These state/parameter values may subject to change on a Slice boundary, and must be provided in each AVC_BSD_OBJECT command. The values set for these variables are retained internally, until they are reset by H/W Asynchronous Reset or changed by the next AVC_BSD_OBJECT command.

Dwords 3-7 are the common in-line data for [DevCTG and DevILK]



| Dword | Bit | Description |
|-------|-------|--|
| 3 | 31 | Concealment Method This field specifies the method used for concealment when error is detected. If set, a copy from collocated macroblock location is performed from the concealment reference indicated by the ConCeal_Pic_Id field. If it is not set, a copy from the current picture is performed using Intra 16x16 Prediction method. 0 – Intra 16x16 Prediction 1 – Inter P Copy |
| | 30 | Init Current_MB_Number When set, the current Slice_Start_MB_Num, Slice_MB_Start_Hor_Pos and Slice_MB_Start_Vert_Pos fields will be used to initialize the Current_MB_Number register. This effectively disables the concealment capability. |
| | 29:28 | Reserved : MBZ |
| | 27:24 | Rewind_Num This field provides the number of MBs or MBAFF pairs to rewind when performing concealment. This is ignored if Slice_Start_MB_Num is smaller than or equal to the Current_MB_Num. Error Concealment MB start position = Current_MB_NUMBER – MIN { Decoded_MB_NUMBER, (MBAFF ? 2 : 1)*Force_Skip_Rewind } Decoded_MB_NUMBER = Number of MBs decoded in the previous slice. |
| | 23:22 | Reserved : MBZ |
| | 21:16 | Conceal_Pic_Id (Concealment Picture ID) This field identifies the picture in the reference list to be used for concealment. This field is only valid if Concealment Method is Inter P Copy. Bit 21 = 0 – Frame Picture = 1 – Field Picture Bit 20:16 – Frame Store Index[4:0] |
| | 15 | Reserved : MBZ |
| | 14 | BSDPrematureComplete Error Handling 1 – Set the interrupt to the driver (provide MMIO registers for MB address R/W) 0 – Ignore the error and continue (masked the interrupt), assume the H/W automatically perform the error handling It occurs in situation where the Slice decode is completed but there are still data in the bitstream. |
| | 13 | Reserved : MBZ |
| | 12 | MPR Error (MV out of range) Handling– what to do when the specific error has occurred 1 – Set the interrupt to the driver (provide MMIO registers for MB address R/W) 0 – Ignore the error and continue (masked the interrupt), assume the H/W automatically perform the error handling |
| | 11 | Reserved : MBZ |
| | 10 | Entropy Error Handling – what to do when the specific error has occurred 1 – Set the interrupt to the driver (provide MMIO registers for MB address R/W) 0 – Ignore the error and continue (masked the interrupt), assume the H/W automatically perform the error handling |

| Dword | Bit | Description |
|-------|-------|--|
| | 9 | Reserved : MBZ |
| | 8 | MB Header Error Handling – what to do when the specific error has occurred 1 – Set the interrupt to the driver (provide MMIO registers for MB address R/W) 0 – Ignore the error and continue (masked the interrupt), assume the H/W automatically perform the error concealment. |
| | 7:4 | Reserved : MBZ |
| | 3:0 | Slice_Type (Slice Type) 0000 – P Slice 0001 – B Slice 0010 – I Slice Other encodings are reserved (note that bits[3:2] must be 0) It is set to the value of the syntax element read from the Slice Header. |
| 4 | 31:30 | Reserved : MBZ |
| | 29:24 | Num_Ref_Idx_L1 (Number of Reference Pictures in Inter-prediction List 1) This field is valid only for decoding a B Slice, for which it is expected to have at least one entry in the reference list L1; otherwise (if Slice Type is not a B Slice), this field must be set to 0. This field can be derived for a B Slice from the Slice Header syntax element NumRefIdxActiveMinus1 as, Num_Ref_Idx_L1 = NumRefIdxActiveMinus1[1] + 1. Format U6 with valid range of [0, 32]. |
| | 23:22 | Reserved : MBZ |
| | 21:16 | Num_Ref_Idx_L0 (Number of Reference Pictures in Inter-prediction List 0) This field is valid for decoding a P or B Slice, for which it is expected to have at least one entry in the reference list L0; otherwise (if Slice Type is not a P or B Slice), this field must be set to 0. This field can be derived for a P or B Slice from the Slice Header syntax element NumRefIdxActiveMinus1 as, Num_Ref_Idx_L0 = NumRefIdxActiveMinus1[0] + 1. Format: U6 with valid range of [0, 32] |
| | 15:11 | Reserved : MBZ |
| | 10:8 | Log2WeightDenomChroma It is the base 2 logarithm of the denominator for all Chroma (Cb and Cr) weighting factors. The value of chroma_log2_weight_denom should be in the range of 0 to 7. It is set to the value of the syntax element read from the Slice Header Pred_Weight_Table(). Format: U3 with valid range of [0, 7]. |
| | 7:3 | Reserved : MBZ |
| | 2:0 | Log2WeightDenomLuma It is the base 2 logarithm of the denominator for all Luma weighting factors. It is set to the value of the syntax element read from the Slice Header Pred_Weight_Table(). Format: U3 with valid range of [0, 7]. |



| Dword | Bit | Description |
|-------|-------|--|
| 5 | 31:30 | <p>Weighted_Pred_Idx (Weighted Prediction Indicator)</p> <p>This field indicates the Weighted Prediction mode for a P or B Slice. It is a combined field corresponding to the syntax element WeightedBiPredIdx or WeightedPredFlag read from the current active PPS.</p> <p>If it is a B-Slice, these 2 bits is interpreted as :</p> <ul style="list-style-type: none"> 00 – specifies the default weighted inter-prediction to be applied 01 – specifies the explicit weighted inter-prediction to be applied 10 – specifies the implicit weighted inter-prediction to be applied 11 – Reserved (not allowed) <p>If it is a P Slice, Weighted_Pred_Idx is interpreted as :</p> <ul style="list-style-type: none"> 00 – disables weighted inter-prediction (default weighted) 01 – enables weighted inter-prediction (explicit weighted) 10-11 – Reserved <p>Only when in B Slice with Weighted_Pred_Idx = 1 (explicit weighted prediction), will there be a L1 and/or a L0 weight+offset tables being sent to the BSD unit through the Slice_State command. Only when in P Slice with Weighted_Pred_Idx = 1, will there be a L0 weight+offset table being sent to the BSD.</p> <p>If Weighted_Pred_Idx != 1 for B Slice or Weighted_Pred_Idx =0 for P Slice, no Slice_State command should be issued to send these tables. If still being issued, the data is read but ignored.</p> |
| | 29 | <p>Direct_Pred_Type (Direct Prediction Type)</p> <p>Type of direct prediction used for B Slices. This field is valid only for Slice_Type = B Slice; otherwise, it must be set to 0.</p> <ul style="list-style-type: none"> 0 – Temporal 1 – Spatial |
| | 28:27 | <p>Disable_Deblocking_Filter_Idx (Disable Deblocking Filter Indicator)</p> <ul style="list-style-type: none"> 00 - filterInternalEdgesFlag is set equal to 1 01 – disable all deblocking operation, no deblocking parameter syntax element is read; filterInternalEdgesFlag is set equal to 0 10 - macroblocks in different slices are considered not available; filterInternalEdgesFlag is set equal to 1 11 – Reserved (not defined in AVC) |
| | 26 | Reserved : MBZ |
| | 25:24 | <p>Cabac_Init_Idx</p> <p>Specifies the index for determining the initialization table used in the context variable initialization process.</p> |
| | 23:22 | Reserved : MBZ |
| | 21:16 | <p>Slice_Qp (Slice Quantization Parameter)</p> <p>Quantization Parameter for current slice. Derived from PPS and slice_delta_qp syntax element in Slice Header.</p> |
| | 15:12 | Reserved : MBZ |



| Dword | Bit | Description |
|-------|-------|---|
| | 11:8 | Slice_Beta_Offset_Div2 Specifies the offset used in accessing the deblocking filter strength tables. It is a sign 2's complement number, ranging from -6 to +6 inclusive. Format: S4 with valid range [-6, 6] |
| | 7:4 | Reserved. MBZ. |
| | 3:0 | Slice_Alpha_C0_Offset_Div2 Specifies the offset used in accessing the deblocking filter strength tables. It is a sign 2's complement number, range from -6 to +6 inclusive. Format: S4 with valid range [-6, 6] |
| 6 | 31:24 | Slice_MB_Start_Vert_Pos (Slice Vertical Position) This field specifies the position in y-direction of the first macroblock in the Slice in unit of macroblocks. |
| | 23:16 | Slice_MB_Start_Hor_Pos (Slice Horizontal Position) This field specifies the position in x-direction of the first macroblock in the Slice in unit of macroblocks. |
| | 15 | Reserved : MBZ |
| | 14:0 | Slice_Start_Mb_Num The MB number (linear MB address in a picture) at the start of a Slice, it must match with the Slice Horizontal Position (Slice_MB_Start_Hor_Pos) and Vertical Position (Slice_MB_Start_Vert_Pos) in the picture. In creating the Phantom Slice for error concealment, this field should set to the total number of MB in the current picture + 1. |
| 7 | 31:8 | Reserved : MBZ |
| | 7 | Fix_Prev_Mb_Skipped Enables an alternative method for decoding mb_skipped, to cope with an encoder that codes a skipped MB as a direct MB with no coefficient. |
| | 6:3 | Reserved : MBZ |
| | 2:0 | First_MB_Bit_Offset (First Macroblock Bit Offset) This field provides the bit offset of the first macroblock of the Slice in the first byte of the input compressed bitstream. This field is the number of valid bits minus 1 of the first byte. For example when this field is set to 7, it indicates that the whole first byte (all 8 bits) are valid; If this field is set to 0, it indicates that only the 1 LSB of the first byte is valid data for the first macroblock. Programming notes: Byte alignment of CABAC. According to AVC Specification, if a slice is coded as CABAC (entropy_coding_flag = CABAC), the first macroblock of a slice is always byte aligned, i.e., the slice header is bit padded to align on byte boundary. So for CABAC case, this field must be set to 7 indicating that the first macroblock is byte aligned. Host software must take care of slice header bit padding to find the first bit of the first macroblock. It should be noted that hardware can take care of the slice header bit padding for most cases, but it falls to deal with one corner condition – when the first byte where the slice header ends is the last byte of a cache line. Format: U3 |



Dwords 8 is always present for [DevILK], and is only present when AES decryption is enabled (Bit31 of DW1 is set to 1) for [DevCTG].

| Dword | Bit | Description |
|-------|------|---|
| 8 | 31:0 | <p>Initial counter value for AES-128 Counter mode decryption. 32-bit initial counter value (which does not need to be secure) for the AES counter mode decryption.</p> <p>This counter value should be applied in the AES Counter mode decryption of the first 16byte chunk of aligned encrypted slice data which contains the 1st starting byte of valid slice bitstream.</p> <p>If Bit stream is AES-128 Counter mode Encrypted is 0, i.e. not encrypted bit stream, this field is ignored and should be MBZ.</p> |

Dwords 9-15 are only supported by [DevILK]

| Dword | Bit | Description |
|-------|-------|---|
| 9 | 31:23 | Reserved. MBZ |
| | 22:0 | Reserved |
| 10 | 31:29 | Packet-based function. Starting bit position of the Bit Vector in the 1st byte of the fetched Bit Vector surface. |
| | 28:16 | Reserved : MBZ |
| | 15:0 | <p>Total Packet Count. This is a total count of all the packets in the packet-based bitstream, including those that are fetched for processing and, also those packets which are skipped (because they belonged to a different stream).</p> <p>Format = U23 in bytes</p> <p>If the bitstream is not packet-based, all the fields in this DWord are ignored and MBZ.</p> |
| 11 | 31:29 | Reserved : MBZ |
| | 28:0 | <p>Packet-based Bitsream Bit Vector Surface starting Byte address. This field specifies the Graphics Memory starting address of the Bit Vector to be used for skipping/fetching bitstream slice packets. This virtual address points to a 4K page and is therefore 4KByte aligned.</p> <p>Each data bit in this surface refers to a bitstream packet and indicates:</p> <p>'0' – The packet is used in the bitstream decode for the current slice, and therefore needs to be fetched.</p> <p>'1' – The packet is not used in the bitstream decode for the current slice, and therefore need not be fetched.</p> <p>If the bitstream is not packet-based, this field is ignored and MBZ.</p> |
| 12 | 31:0 | <p>Weight128LumaL0 – One bit for each of the I0/I1 entries for luma and chroma components. This bit should be set by driver whenever the 16 bit weight value from the application is equal to 128. this will be passed on as inline data in the media object command.</p> <p>Each of the bits is for one entry in the L0 list with the LSB representing if index 0 is 128 or not for the luma component</p> |



| Dword | Bit | Description |
|-------|------|--|
| 13 | 31:0 | <p>Weight128LumaL1 – One bit for each of the I0/I1 entries for luma and chroma components. This bit should be set by driver whenever the 16 bit weight value from the application is equal to 128. this will be passed on as inline data in the media object command.</p> <p>Each of the bits is for one entry in the L1 list with the LSB representing if index 0 is 128 or not for the luma component.</p> |
| 14 | 31:0 | <p>Weight128ChromaL0 – One bit for each of the I0/I1 entries for luma and chroma components. This bit should be set by driver whenever the 16 bit weight value from the application is equal to 128. this will be passed on as inline data in the media object command.</p> <p>Each of the bits is for one entry in the L0 list with the LSB representing if index 0 is 128 or not for the chroma component</p> |
| 15 | 31:0 | <p>Weight128ChromaL1 – One bit for each of the I0/I1 entries for luma and chroma components. This bit should be set by driver whenever the 16 bit weight value from the application is equal to 128. this will be passed on as inline data in the media object command.</p> <p>Each of the bits is for one entry in the L1 list with the LSB representing if index 0 is 128 or not for the chroma component.</p> |

2.7.3.2 Indirect Data Format

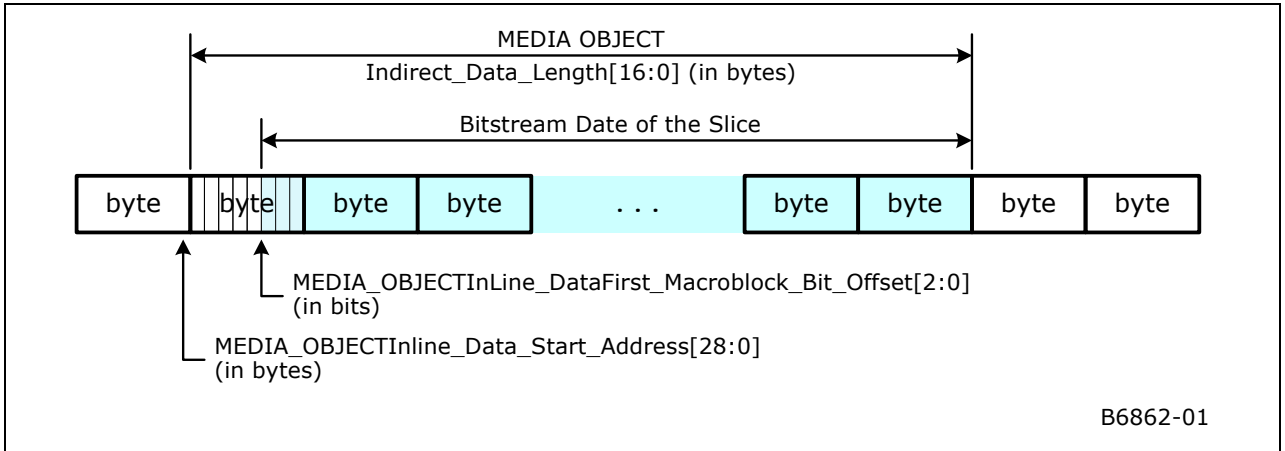
Each AVC_BSD_OBJECT command corresponds to the processing of one slice of a picture (a picture can have only one slice). All syntax and derived parameters above Slice Data Layer are passed in either using State commands or as Inline data of the AVC_BSD_OBJECT command, and the Slice Data Layer and below are passed in as indirect data.

The indirect data start address in AVC_BSD_OBJECT specifies the starting Graphics Memory address of the bitstream data that follows the slice header. It provides the byte address for the first macroblock of the slice. Together with the First MB Bit Offset field, it provides the starting bit location of the MB within the compressed bitstream.

The indirect data length in AVC_BSD_OBJECT provides the length in bytes of the bitstream data for this slice. It includes the first byte of the first MB (includes the one that contains the FirstMBBiteOffset) and the last non-zero byte of the last MB in the slice. H/W ignores the contents after the last non-zero byte.



Figure 2-1. Indirect data buffer for a slice



2.8 AVC_BSD Output Definitions

Three data buffers are output from the BSD engine:

- AVC-IT Command Buffer – MB info signals from AC to AM (mb_avail, cbp) + AM (avail, cbp)
- AVC-IT Data Buffer – IPCM or Coeff + MV/Weight+Offset
- AVC-ILDB Data Buffer

The format of the AVC-IT Command Buffer can be found in the *Media* Chapter, under section *MEDIA_OBJECT_EX Command* as well as its subsection *In-line Data Format in AVC-IT Mode*.

The format of the AVC-IT Data Buffer can be found in the *Media* Chapter, under section *MEDIA_OBJECT_EX Command* and subsection *Indirect Data Format in AVC-IT Mode*.

The format of the AVC-LF Data Buffer can be found in the *Data Port* Chapter, under section *AVC Loop Filter Read*.



2.8.1 AVC-IT Command Buffer Description

2.8.2 AVC-IT Data Buffer Description

2.8.2.1 Motion Vectors

Motion vectors (x and y components) are sent to the MC as indirect data in memory together with residual data. In non-I_PCM mode, they are the packed computed MVs (Mv_{predicted} + MVDelta) of a MB; in I_PCM mode, there is no MV information need to be sent. MV can be generated for DMV-spatial, DMV-temporal or motion-search MV. The two components of MV (MV_x, MV_y) are signed extended into a 16-bit signed number each and are packed into a single Dword. MVs of a MB are packed in 512-bit cache lines (or in half-cacheline ???). Based on the MB_type, we can determine the number of MVs being present in the corresponding MB. There can be up to 2, 4, 8, 16 and 32 MVs for a MB in both directions. A MV can associate with a 4x4,4x8,8x4 sub-block, a 8x8, 8x16, 16x8 block or 16x16 MB, and for a direction (L0 and L1). The MV for sub-blocks are sent in a raster order within a block, and in raster order in respect to the blocks within a MB. Zero-value MVs are sent as is. When L1 is not in use, no MV are sent for L1.

| DWord | Bit | Description |
|----------------|-------|---|
| For each Dword | 31:16 | MV_y Signed extended 2's complement |
| | 15:0 | MV_x Signed extended 2's complement |

2.8.2.1.1 MV Organization for 16x16 Block Size

The MVs are specified in a fixed data structure, which is organized as if there are always four 8x8 blocks. Hence, the real MV(s) is (are) replicated across the remaining entries.

2.8.2.1.2 MV Organization for 16x8 Block

A MB that is 16x8 partitioned consists of 2 parts – top and bottom partitions. However, its MVs are specified as if the MB is broken down into four 8x8 blocks. That is, each 16x8 partition is treated as two identical 8x8 blocks.

2.8.2.1.3 MV Organization for 8x16 Block

A MB that is 8x16 partitioned consists of 2 parts – left and right partitions. However, its MVs are specified as if the MB is broken down into four 8x8 blocks. That is, each 8x16 partition is treated as two identical 8x8 blocks. The MV are specified as if the MB is broken down into four 8x8 blocks.

2.8.2.1.4 MV Organization for 8x8 Block Without Subpartition

A MB that is 8x8 partitioned consists of 4 parts – top-left (0), top-right(1), bottom-left(2) and bottom-right(3) partitions. The MVs are specified for each 8x8 blocks.



2.8.2.1.5 MV Organization for 8x8 Block with Subpartition

A MB that has at least one sub-partitioning (4x4, 8x4 or 4x8), the entire MB is broken down into 16 4x4 units for the specification of MV. That is, the MVs are specified for each 4x4 blocks.

The 4x4 block are scanned in the following order :

| | | | |
|----|----|----|----|
| 0 | 1 | 4 | 5 |
| 2 | 3 | 6 | 7 |
| 8 | 9 | 12 | 13 |
| 10 | 11 | 14 | 15 |

Replication Rules:

- For an 8x8 partition, the 8x8 MVs are replicated into all the four 4x4 blocks.
- For an 4x8 subpartition within an 8x8 partition, each 4x8 is broken down into 2 4x4 stacking vertically. The 4x8 MVs are replicated into both 4x4 blocks.
- For an 8x4 subpartition within an 8x8 partition, each 8x4 is broken down into 2 4x4 stacking horizontally. The 8x4 MVs are replicated into both 4x4 blocks.
- For a 4x4 subpartition within an 8x8 partition, each 4x4 has its own MVs.
- For L0 MV
 - If L0 interprediction exists, the corresponding L0 MV is used
 - Else if L1 interprediction exists (of the same block), set to the same as L1 MV
 - Else set to 0
- For L1 MV
 - If L1 interprediction exists, the corresponding L1 MV is used
 - Else if L0 interprediction exists (of the same block), set to the same as L0 MV
 - Else set to 0

MBLK data (containing the packed MV/Weight,Offset/Coeff) offset in the indirect buffer can be dword aligned. MV/Weight,offset/Coeff are packed in memory as dword aligned data structure (constraint to Dword align for a buffer), VFE has to derive the size for each based on bMSize. BSD output is half_cacheline aligned.

Block address (right to left)

MV write out first, then coeff



Weight offset does not go down to 4x4.

Replicate Weight/Offset, decomp 16x8 to 8x8, with the same Weight/Offset.

2.8.2.2 Inter-prediction Weight and Offset Block

Weight and offset are packed together next to each other. A weight and an offset are each 1 byte. The high byte is always the weight and the lower byte is always the offset. The smallest block size for Weight and Offset specifications is 8x8 (no subpartitioning, since interprediction mode does not go below 8x8).

2.8.2.2.1 Weight+Offset Organization for 16x16 Block Size

The Weight+Offset are specified in a fixed data structure. Hence, the real Weight+Offsets are replicated across the remaining entries.

| | DWord | Description |
|--|----------------|---|
| | DW0 [15:0] | L0 Weight+Offset_Y0 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Y0 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Y0 Else set to 0 |
| | DW0 [31:16] | L1 Weight+Offset_Y0 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Y0 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Y0 Else set to 0 |
| | DW1 [15:0] | L0 Weight+Offset_Cb0 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Cb0 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Cb0 Else set to 0 |
| | DW1 [31:16] | L1 Weight+Offset_Cb0 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Cb0 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Cb0 Else set to 0 |
| | DW2 [15:0] | L0 Weight+Offset_Cr0 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Cr0 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Cr0 Else set to 0 |
| | DW2 [31:16] | L1 Weight+Offset_Cr0 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Cr0 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Cr0 Else set to 0 |



| | DWord | Description |
|--|-------|--|
| | DW3 | Reserved. |
| | DW4 | Replicated DW0 (L0 Weight+Offset_Y0 and L1 Weight+Offset_Y0) |
| | DW5 | Replicated DW1 (L0 Weight+Offset_Cb0 and L1 Weight+Offset_Cb0) |
| | DW6 | Replicated DW2 (L0 Weight+Offset_Cr0 and L1 Weight+Offset_Cr0) |
| | DW7 | Reserved. |

2.8.2.2.2 Weight+Offset Organization for non-16x16 Block Size

The Weight+Offset are specified in a fixed data structure. It is organized as if the MB is broken into four 8x8 blocks. Hence, the real Weight+Offsets are replicated across the remaining entries.

Replication Rules:

- For 8x16 partition, each 8x16 is broken down into 2 8x8 stacking vertically. The 8x16 Weight+Offset are replicated into the two 8x8 blocks.
- For 16x8 partition, each 16x8 is broken down into 2 8x8 stacking horizontally. The 16x8 Weight+Offset are replicated into the two 8x8 blocks.
- For 8x8 partition, the MB is broken down into 4 8x8 blocks. Each has its own set of Weight and Offset.

| | DWord | Description |
|--|----------------|---|
| | DW0 [15:0] | L0 Weight+Offset_Y0 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Y0 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Y0 Else set to 0 |
| | DW0 [31:16] | L1 Weight+Offset_Y0 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Y0 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Y0 Else set to 0 |
| | DW1 [15:0] | L0 Weight+Offset_Cb0 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Cb0 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Cb0 Else set to 0 |



| | DWord | Description |
|--|----------------|---|
| | DW1 [31:16] | L1 Weight+Offset_Cb0 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Cb0 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Cb0 Else set to 0 |
| | DW2 [15:0] | L0 Weight+Offset_Cr0 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Cr0 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Cr0 Else set to 0 |
| | DW2 [31:16] | L1 Weight+Offset_Cr0 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Cr0 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Cr0 Else set to 0 |
| | DW3 | Reserved. |
| | DW4 [15:0] | L0 Weight+Offset_Y1 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Y1 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Y1 Else set to 0 |
| | DW4 [31:16] | L1 Weight+Offset_Y1 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Y1 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Y1 Else set to 0 |
| | DW5 [15:0] | L0 Weight+Offset_Cb1 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Cb1 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Cb1 Else set to 0 |
| | DW5 [31:16] | L1 Weight+Offset_Cb1 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Cb1 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Cb1 Else set to 0 |
| | DW6 [15:0] | L0 Weight+Offset_Cr1 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Cr1 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Cr1 Else set to 0 |



| | DWord | Description |
|--|-----------------|--|
| | DW6 [31:16] | L1 Weight+Offset_Cr1 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Cr1 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Cr1 Else set to 0 |
| | DW7 | Reserved. |
| | DW8 [15:0] | L0 Weight+Offset_Y2 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Y2 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Y2 Else set to 0 |
| | DW8 [31:16] | L1 Weight+Offset_Y2 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Y2 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Y2 Else set to 0 |
| | DW9 [15:0] | L0 Weight+Offset_Cb2 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Cb2 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Cb2 Else set to 0 |
| | DW9 [31:16] | L1 Weight+Offset_Cb2 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Cb2 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Cb2 Else set to 0 |
| | DW10 [15:0] | L0 Weight+Offset_Cr2 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Cr2 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Cr2 Else set to 0 |
| | DW10 [31:16] | L1 Weight+Offset_Cr2 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Cr2 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Cr2 Else set to 0 |
| | DW11 | Reserved. |



| | DWord | Description |
|--|-----------------|--|
| | DW12 [15:0] | L0 Weight+Offset_Y3 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Y3 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Y3 Else set to 0 |
| | DW12 [31:16] | L1 Weight+Offset_Y3 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Y3 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Y3 Else set to 0 |
| | DW13 [15:0] | L0 Weight+Offset_Cb3 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Cb3 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Cb3 Else set to 0 |
| | DW13 [31:16] | L1 Weight+Offset_Cb3 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Cb3 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Cb3 Else set to 0 |
| | DW14 [15:0] | L0 Weight+Offset_Cr3 The high byte is weight and the low byte is offset If L0 interprediction exists, it is equal to L0 Weight+Offset_Cr3 Else if L1 interprediction exists, it is a replication of L1 Weight+Offset_Cr3 Else set to 0 |
| | DW14 [31:16] | L1 Weight+Offset_Cr3 The high byte is weight and the low byte is offset If L1 interprediction exists, it is equal to L1 Weight+Offset_Cr3 Else if L0 interprediction exists, it is a replication of L0 Weight+Offset_Cr3 Else set to 0 |
| | DW15 | Reserved. |

2.8.2.3 Quantized DCT Coefficient Block

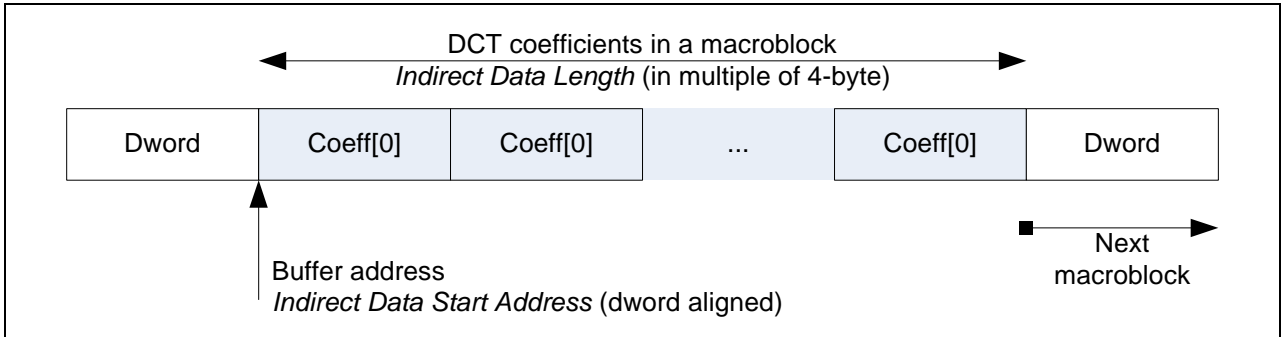
In non-I_PCM mode, quantized DCT coefficients are sent to AVC-IT as indirect data in memory (AVC-IT Data Buffer). They are organized on a per MB basis (does it include skip MB and are they zero filled), i.e. all coefficients of a MB are packed together; in I_PCM mode, the AVC-IT buffer is packed with actual 8-bit pixel samples on a per MB basis.

Coefficients are packed in 512-bit cache lines (Note: should be dword aligned, not cache line aligned???)



Only the non-zero quantized DCT coefficients within a MB are sent and they are packed in the order of 4x4DCY (if any), 8x8Y0, 8x8Y1, 8x8Y2, 8x8Y3, 2x2DCCb (if any), 2x2DCCr (if any), 8x8Cb4 and 8x8Cr5, as shown in **Error! Reference source not found..** For I16x16Intra mode, the DCY is the very first block; in all other prediction modes, 8x8Y0 is the first block.

Structure of the Quantized DCT Coefficients Data Buffer



Each non-zero coefficient in the indirect data buffer is contained in a doubleword-size (32-bit) data structure containing the coefficient index, end of block (EOB) flag and the fixed-point coefficient value in 2's complement form.

As shown in **Error! Reference source not found..**, *index* is the row major 'raster' index of the coefficient within an 8x8 block. DCT coefficient is a 16-bit value in 2's complement, which is clamped to a 12-bit signed value by the host.

Structure of a quantized DCT coefficient block

| DWord | Bit | Description |
|----------------|-------|---|
| For each Dword | 31:16 | Quantized DCT Coefficient Value or Pixel sample value This field contains the value of the non-zero quantized DCT coefficient in 2's complement. |
| | 15:7 | Reserved: MBZ |
| | 6:1 | Index This field specifies the raster-scan address (raw address) of the quantized DCT coefficient (or pixel sample in I_PCM) within their corresponding block (some of size 8x8, and some are 4x4 or 2x2) . For example, coefficient in an 8x8Y0 at location (row, column) = (0, 0) has an index of 0; that at (2, 3) has an index of $2*8 + 3 = 19$. For 4x4 YDC block, index ranges from 0 to 15; for 2x2 CrDC/CbDC, ranges from 0 to 3. Format = U6 Range = [0, 63] |
| | 0 | EOB (End of Block) This field indicates whether the quantized DCT coefficient is the last one of the current block (8x8, 4x4 or 2x2). |

2.8.2.4 I_PCM Pixel Sample Block



| DWord | Bit | Description |
|----------------|------|--|
| For each Dword | 31:0 | 8-bit Packed Pixel sample value Each Dword contains 4 pixel sample values (zero or not) packed together. |

2.8.2.5 Motion Prediction Reference Indices

Motion prediction reference indices for L0 and L1 are sent separately to the MC as inline data in memory. In non-I_PCM mode, they are the packed for a MB and are aligned with the associated MVs; in I_PCM mode, there is no MV information need to be sent. MV can be generated for DMV-spatial, DMV-temporal or motion-search MV. Each reference index is a 4bit unsigned number, except that 1111 (-1) is used to mark for a non-existing reference picture. 8 reference indices can be packed into a single Dword. Reference indices of a MB are packed in 512-bit cache lines (or in half-cacheline ???). Based on the MB_type, we can determine the number of reference indices being present in the corresponding MB. A Reference indices can associate with a a 8x8, 8x16, 16x8 block or 16x16 MB, and for a direction (L0 and L1). The Reference indices for sub-blocks are sent in a raster order within a block, and in raster order in respect to the blocks within a MB. When L1 is not in use, no reference indices are sent for L1.

2.8.2.6 MB Neighbor Availability in Intra-Prediction Modes

mb_nb_avail_intra[4:0] provides the individual MB availability information as follows. [4:3] bit ordering is under discussion. Current MB is labelled as X. For a MB pair, both the top and the bottom have the same picture type.

- mb_nb_avail_intra[0] : MB D (top left neighbor of current MB X) availability
- mb_nb_avail_intra[1] : MB C (top right neighbor of current MB X) availability
- mb_nb_avail_intra[2] : MB B (top neighbor of current MB X) availability
- mb_nb_avail_intra[3] : MB A availability. A is the left neighbor of the current MB X, or A is the top MB of the left neighbor of the current MB pair X.
- mb_nb_avail_intra[4] : MB E availability. E is the bottom MB of the left neighbor of the current MB pair X.

mb_nb_avail_intra[4 :0] is derived from the internal intermediate variable mb_avail_i [0..7] (internally generated by AM unit based on mb_nb_avail[0..3]), which in turns is derived from the initial variable mb_nb_avail[3:0].

mb_nb_avail[3:0] is derived for a MB (a MB unit) if in non-MBAFF mode or for a MB pair (a MB unit) if in MBAFF mode, from the following conditions :

1. A MB is marked as “Not Available” if anyone of the following conditions is true :
 - (MbAddr – Start_MB_addr) < 0 (error condition), for each Slice
 - MbAddr > CurrMBAddr (go beyond the current MB location, i.e. to the right or below)
 - The MB with address MbAddr belongs to a different slice than the MB with address CurrMbAddr
 - Set the mb_nb_avail[i] accordingly for the MB; mb_nb_avail[i] = 0 if “not available”, i = 0 to 3
2. For each non-pair MB with address CurrMbAddr, check the following availability, and set the mb_nb_avail[i] accordingly for the corresponding MB
 - mbAddrA = CurrMbAddr – 1,



- mbAddrA is marked as not available when $\text{CurrMbAddr} \% \text{PicWidthInMbs}$ is equal to 0 (i.e. the MB with address CurrMbAddr is located along the left edge of the picture)
- $\text{mbAddrB} = \text{CurrMbAddr} - \text{PicWidthInMbs}$
- $\text{mbAddrC} = \text{CurrMbAddr} - \text{PicWidthInMbs} + 1$
 - mbAddrC is marked as not available when $(\text{CurrMbAddr} + 1) \% \text{PicWidthInMbs}$ is equal to 0 (i.e. the MB with address CurrMbAddr is located along the right edge of the picture)
- $\text{mbAddrD} = \text{CurrMbAddr} - \text{PicWidthInMbs} - 1$
 - mbAddrD is marked as not available when $\text{CurrMbAddr} \% \text{PicWidthInMbs}$ is equal to 0 (i.e. the MB with address CurrMbAddr is located along the left edge of the picture)
 -

| | | |
|---|--|--|
| mbAddrD D (top-left) mb_nb_avail[0] mb_avail_i[0,1] | mbAddrB B (top) mb_nb_avail[2] mb_avail_i[4,5] | mbAddrC C (top-right) mb_nb_avail[1] mb_avail_i[2,3] |
| mbAddrA A (left) mb_nb_avail[3] mb_avail_i[6,7] | X CurrMbAddrX | N/A |
| N/A | N/A | N/A |

3. For each MB pair (in MBAFF mode) with address CurrMbAddr, check the following availability, and set the mb_nb_avail[i] accordingly for the corresponding neighbor MB pair (Note that in MBAFF mode, the top MB always has an even address and the bottom MB has an odd address; also in MBAFF, both MBs must have the same availability value, hence no need to specify top or bottom MB in the mb_nb_avail[i]. However, the top and the bottom MB may code with different modes - intra or inter).
 - $\text{mbAddrA} = 2 * (\text{CurrMbAddr} / 2 - 1)$
 - mbAddrA is marked as not available when $(\text{CurrMbAddr} / 2) \% \text{PicWidthInMbs}$ is equal to 0 (i.e. the MB pair with address CurrMbAddr is located along the left edge of the picture)
 - $\text{mbAddrB} = 2 * (\text{CurrMbAddr} / 2 - \text{PicWidthInMbs})$
 - $\text{mbAddrC} = 2 * (\text{CurrMbAddr} / 2 - \text{PicWidthInMbs} + 1)$
 - mbAddrC is marked as not available when $(\text{CurrMbAddr} / 2 + 1) \% \text{PicWidthInMbs}$ is equal to 0 (i.e. the MB with address CurrMbAddr is located along the right edge of the picture)
 - $\text{mbAddrD} = 2 * (\text{CurrMbAddr} / 2 - \text{PicWidthInMbs} - 1)$
 - mbAddrD is marked as not available when $(\text{CurrMbAddr} / 2) \% \text{PicWidthInMbs}$ is equal to 0 (i.e. the MB with address CurrMbAddr is located along the left edge of the picture)
 - $\text{mbAddrE} = \text{mbAddrA} + 1$
 - its MB availability should be the same as that of mbAddrA

| | | |
|---|--|--|
| mbAddrD D0 mb_nb_avail[0] mb_avail_i[0] | mbAddrB B0 mb_nb_avail[2] mb_avail_i[4] | mbAddrC C0 mb_nb_avail[1] mb_avail_i[2] |
| mbAddrD+1 D1 mb_nb_avail[0] mb_avail_i[1] | mbAddrB+1 B1 mb_nb_avail[2] mb_avail_i[5] | mbAddrC+1 C1 mb_nb_avail[1] mb_avail_i[3] |
| mbAddrA A0 mb_nb_avail[3] mb_avail_i[6] | CurrMbAddrX X0 or | N/A |
| mbAddrA+1 A1/E (mbAddrE) mb_nb_avail[3] mb_avail_i[7] | CurrMbAddrX X1 | N/A |

Additional conditions for each Intra_prediction Modes can further refine the MB neighbor's (N = A, B, C, D, or E) availability:

- Intra_4x4 prediction
 - If any of the following conditions is true, the neighbor is marked as “not available for Intra_4x4 prediction”
 - mbAddrN is not available,
 - the macroblock mbAddrN is coded in Inter prediction mode and constrained_intra_pred_flag is equal to 1.
 - the top right 4x4 block (luma4x4BlkIdx = 4 or 12) to the block with luma4x4BlkIdx equal to 3 or 11 (because not decoded yet) – Kernel will take care of it for internal blocks.
- Intra_8x8 prediction
 - If any of the following conditions is true, the neighbor is marked as “not available for Intra_8x8 prediction”
 - mbAddrN is not available,
 - the macroblock mbAddrN is coded in Inter prediction mode and constrained_intra_pred_flag is equal to 1,



- Intra_16x16 prediction
 - If any of the following conditions is true, the neighbor is marked as “not available for Intra_16x16 prediction”
 - mbAddrN is not available,
 - the macroblock mbAddrN is coded in Inter prediction mode and constrained_intra_pred_flag is equal to 1.

The above derivation of MB availability flag can be summarized into one equation:

$mb_avail_i[i] = mb_nb_avail[i/2] \& (is_intra_MB[i] | ! constrained_intra_pred_flag), i = 0 \text{ to } 7.$
mb_avail_i[0] – for MB D0
mb_avail_i[1] – for MB D1
mb_avail_i[2] – for MB C0
mb_avail_i[3] – for MB C1
mb_avail_i[4] – for MB B0
mb_avail_i[5] – for MB B1
mb_avail_i[6] – for MB A0
mb_avail_i[7] – for MB A1

is_intra_MB[1:0] - Upleft_mb_iblocks[0:1], i.e. D1 and D0

- Specifies if the upper left MB or the upper left MB pair to the current MB is an intra MB or not
- For the upper left neighbor being a MB pair, the Upleft_mb_iblocks[1] is for the Top MB and the Upleft_mb_iblocks[0] is for the Bottom MB

is_intra_MB[3:2] - Upright_mb_iblocks[0:1], i.e. C1 and C0

- Specifies if the upper right MB or the upper right MB pair to the current MB is an intra MB or not
- For the upper right neighbor being a MB pair, the Upright_mb_iblocks[1] is for the Top MB and the Upright_mb_iblocks[0] is for the Bottom MB

is_intra_MB[5:4] - Up_mb_iblocks[0:1], i.e. B1 and B0

- Specifies if the upper MB or the upper MB pair to the current MB is an intra MB or not
- For the upper neighbor being a MB pair, the Up_mb_iblocks[1] is for the Top MB and the Up_mb_iblocks[0] is for the Bottom MB

is_intra_MB[7:6] - left_mb_iblocks[0:1], i.e. A1 and A0

- Specifies if the left MB or the left MB pair to the current MB is an intra MB or not
- For the left neighbor being a MB pair, the Left_mb_iblocks[1] is for the Top MB and the Left_mb_iblocks[0] is for the Bottom MB

To speed up the MC kernels performance for Intra-Prediction modes, the 8-bit MB availability flags Mb_avail_i[7:0] are refined to reduce the number of kernel instructions. The modified availability information has five bits, mb_nb_avail_intra[4:0]. Bit mb_nb_avail_intra[4] is used only in MBAFF mode. The only situation when



`mb_nb_avail_intra[4]` and `mb_nb_avail_intra[3]` may be different is when the current MB pair is a field MB and the left neighbor MB pair is a frame MB.

The diagram shows the physical storage location for a MB pair (16 pixel x 32 rows, numbered from 0 to 31). Logically, the top MB (16x16) of a MB pair contains all the even lines (labelled with index 0, e.g. A0), and the bottom MB (16x16) of a MB pair contains all the odd lines (labelled with index 1, e.g., A1).



| | | | | |
|----|----|-----------------|----|--------------|
| 0 | A0 | Field MB | 0 | |
| 1 | A1 | | 1 | |
| 2 | A0 | Top MB of | 2 | |
| 3 | A1 | | 3 | |
| 4 | A0 | Left Neighbor A | 4 | |
| 5 | A1 | | 5 | |
| 6 | A0 | | 6 | |
| 7 | A1 | | 7 | Frame MB |
| 8 | A0 | | 8 | Top MB of |
| 9 | A1 | | 9 | Current MB X |
| 10 | A0 | | 10 | X0 |
| 11 | A1 | Field MB | 11 | |
| 12 | A0 | | 12 | |
| 13 | A1 | Bottom MB of | 13 | |
| 14 | A0 | | 14 | |
| 15 | A1 | Left Neighbor A | 15 | |
| 16 | A0 | | 16 | |
| 17 | A1 | | 17 | |
| 18 | A0 | | 18 | Frame MB |
| 19 | A1 | | 19 | Bottom MB of |
| 20 | A0 | | 20 | Current MB X |
| 21 | A1 | | 21 | X1 |
| 22 | A0 | | 22 | |
| 23 | A1 | | 23 | |
| 24 | A0 | | 24 | |
| 25 | A1 | | 25 | |
| 26 | A0 | | 26 | |
| 27 | A1 | | 27 | |
| 28 | A0 | | 28 | |
| 29 | A1 | | 29 | |
| 30 | A0 | | 30 | |
| 31 | A1 | | 31 | |



| | | | |
|----|---------------|-------|--------------|
| 0 | | 0 X0 | Field MB |
| 1 | | 1 X1 | |
| 2 | | 2 X0 | Top MB of |
| 3 | | 3 X1 | |
| 4 | | 4 X0 | Current MB X |
| 5 | | 5 X1 | |
| 6 | | 6 X0 | |
| 7 | Frame MB | 7 X1 | |
| 8 | Top MB of | 8 X0 | |
| 9 | Left Neighbor | 9 X1 | |
| 10 | A0 | 10 X0 | |
| 11 | | 11 X1 | Field MB |
| 12 | | 12 X0 | |
| 13 | | 13 X1 | Bottom MB of |
| 14 | | 14 X0 | |
| 15 | | 15 X1 | Current MB X |
| 16 | | 16 X0 | |
| 17 | | 17 X1 | |
| 18 | Frame MB | 18 X0 | |
| 19 | Bottom MB of | 19 X1 | |
| 20 | Left Neighbor | 20 X0 | |
| 21 | A1 | 21 X1 | |
| 22 | | 22 X0 | |
| 23 | | 23 X1 | |
| 24 | | 24 X0 | |
| 25 | | 25 X1 | |
| 26 | | 26 X0 | |
| 27 | | 27 X1 | |
| 28 | | 28 X0 | |
| 29 | | 29 X1 | |
| 30 | | 30 X0 | |
| 31 | | 31 X1 | |



| | | | | | |
|----|----|-----------------|----|----|--------------|
| 0 | A0 | Field MB | 0 | X0 | Field MB |
| 1 | A1 | | 1 | X1 | |
| 2 | A0 | Top MB of | 2 | X0 | Top MB of |
| 3 | A1 | | 3 | X1 | |
| 4 | A0 | Left Neighbor A | 4 | X0 | Current MB X |
| 5 | A1 | | 5 | X1 | |
| 6 | A0 | | 6 | X0 | |
| 7 | A1 | | 7 | X1 | |
| 8 | A0 | | 8 | X0 | |
| 9 | A1 | | 9 | X1 | |
| 10 | A0 | | 10 | X0 | |
| 11 | A1 | Field MB | 11 | X1 | Field MB |
| 12 | A0 | | 12 | X0 | |
| 13 | A1 | Bottom MB of | 13 | X1 | Bottom MB of |
| 14 | A0 | | 14 | X0 | |
| 15 | A1 | Left Neighbor A | 15 | X1 | Current MB X |
| 16 | A0 | | 16 | X0 | |
| 17 | A1 | | 17 | X1 | |
| 18 | A0 | | 18 | X0 | |
| 19 | A1 | | 19 | X1 | |
| 20 | A0 | | 20 | X0 | |
| 21 | A1 | | 21 | X1 | |
| 22 | A0 | | 22 | X0 | |
| 23 | A1 | | 23 | X1 | |
| 24 | A0 | | 24 | X0 | |
| 25 | A1 | | 25 | X1 | |
| 26 | A0 | | 26 | X0 | |
| 27 | A1 | | 27 | X1 | |
| 28 | A0 | | 28 | X0 | |
| 29 | A1 | | 29 | X1 | |
| 30 | A0 | | 30 | X0 | |
| 31 | A1 | | 31 | X1 | |



| | | | |
|----|---------------|----|--------------|
| 0 | | 0 | |
| 1 | | 1 | |
| 2 | | 2 | |
| 3 | | 3 | |
| 4 | | 4 | |
| 5 | | 5 | |
| 6 | | 6 | |
| 7 | Frame MB | 7 | Frame MB |
| 8 | Top MB of | 8 | Top MB of |
| 9 | Left Neighbor | 9 | Current MB X |
| 10 | A0 | 10 | X0 |
| 11 | | 11 | |
| 12 | | 12 | |
| 13 | | 13 | |
| 14 | | 14 | |
| 15 | | 15 | |
| 16 | | 16 | |
| 17 | | 17 | |
| 18 | Frame MB | 18 | Frame MB |
| 19 | Bottom MB of | 19 | Bottom MB of |
| 20 | Left Neighbor | 20 | Current MB X |
| 21 | A1 | 21 | X1 |
| 22 | | 22 | |
| 23 | | 23 | |
| 24 | | 24 | |
| 25 | | 25 | |
| 26 | | 26 | |
| 27 | | 27 | |
| 28 | | 28 | |
| 29 | | 29 | |
| 30 | | 30 | |
| 31 | | 31 | |

| D0 | B0 | C0 |
|-----------------------------|---------------------|------------------|
| D1 (Frame/Field MB) | B1 (Frame/Field MB) | C1 (Frame/Field) |
| A = E = A0 (Frame/Field MB) | X0 Frame MB | |
| A = E = A1 (Field MB) | X1 Frame MB | |
| A = E = A0 (Frame MB) | | |

| | | |
|-------------------------|---------------|---------------|
| D0 (Field MB) | B0 (Field MB) | C0 (Field MB) |
| D1 (Frame MB) | B1 (Frame MB) | C1 (Frame MB) |
| A = A0 (Frame/Field MB) | X0 Field MB | |
| E = A1 (Frame MB) | X1 Field MB | |
| E = A0 (Field MB) | | |

| D0 | B0 | C0 |
|--------------------------|---------------|----------------------|
| D1 | B1 | C1 |
| D = A0 (Frame/Field MB) | X0 (Frame MB) | N/A (Frame/Field MB) |
| A = E = A0 (Field MB) | | |
| D = A1 (Field MB) | X1 Frame MB | |
| A = E = A1 (Frame/Field) | | |



| D0 | B0 | C0 |
|---|------------------|------------------|
| D1 (Frame/Field MB) | B1 (Frame/Field) | C1 (Frame/Field) |
| A = A0 (Frame MB) A = A1 (Field MB) E = A1 (Frame/Field MB) | X0 Field MB | |
| | X1 Field MB | |

The Mb_nb_avail_intra[4:0] is derived from the following information :

- The current MB decoding is in MBAFF or not
- current_mb_y[0] (=0 for even address, i.e. the top MB; =1 for odd address, i.e. bottom MB of a pair)
 - It is equivalent to current_mb_is_bottom flag
- mb_field_flag of the current MB (0 is a frame MB, 1 is a field MB)
- current_mbu_top_avail
 - The availability of the top MB of the current MB unit (a single MB if non-MBAFF, a MB pair if MBAFF). It is relevant only when the current MB is the bottom MB
 - It is equal to (last_mb_i | ! constrained_intra_pred_flag)
- mb_nb_field_flags[3:0] (the field flags for the 4 neighbor MBs)
 - mb_nb_field_flags [0] ; left neighbor MB; labelled as A
 - mb_nb_field_flags [1] : top neighbor MB; labelled as B
 - mb_nb_field_flags [2] : top right neighbor MB; labelled as C
 - mb_nb_field_flags [3] : top left neighbor MB; labelled as D
- The mb_avail_i[7:0] for the current MB



AVC Spec from Portion of the Table- Specification of mbAddrN and yM

| | | | | | | | | |
|-------|---------------------------|---|---------|---------|----------------|-----------------------------|-------------|--------------------------------|
| < 0 | $0 \dots \text{maxH} - 1$ | 1 | mbAddrA | 1 | | mbAddrA | yN | |
| | | | | 0 | $yN \% 2 == 0$ | mbAddrA | $yN \gg 1$ | |
| | | | | 0 | $yN \% 2 != 0$ | mbAddrA + 1 | $yN \gg 1$ | |
| | | | 0 | mbAddrA | 1 | | mbAddrA + 1 | yN |
| | | | | | 0 | $yN \% 2 == 0$ | mbAddrA | $(yN + \text{maxH}) \gg 1$ |
| | | | | | 0 | $yN \% 2 != 0$ | mbAddrA + 1 | $(yN + \text{maxH}) \gg 1$ |
| | | 0 | 1 | mbAddrA | 1 | $yN < (\text{maxH} / 2)$ | mbAddrA | $yN \ll 1$ |
| | | | | | 0 | $yN \geq (\text{maxH} / 2)$ | mbAddrA + 1 | $(yN \ll 1) - \text{maxH}$ |
| | | | | | 0 | | mbAddrA | yN |
| | | | 0 | mbAddrA | 1 | $yN < (\text{maxH} / 2)$ | mbAddrA | $(yN \ll 1) + 1$ |
| | | | | | 1 | $yN \geq (\text{maxH} / 2)$ | mbAddrA + 1 | $(yN \ll 1) + 1 - \text{maxH}$ |
| | | | | | 0 | | mbAddrA + 1 | yN |

The AVC table above is translated to the column A and E below. The rest of the table is derived accordingly.

$\text{Mb_avail_i}[7:0] = \text{A1 (or E), A0, B1, B0, C1, C0, D1, D0}$

| | | D | | B | | C | | A | | E | |
|-------------------|---------|---------|---------|---------|---------|---------|---------|---------|--------------|---------|--------------|
| | | A-Frame | A-Field | A-Frame | A-Field | A-Frame | A-Field | A-Frame | A-Field | A-Frame | A-Field |
| X_0 (Top) | X-Frame | D_1 | D_1 | B_1 | B_1 | C_1 | C_1 | A_0 | $A_0 \& A_1$ | A_0 | $A_0 \& A_1$ |
| | X-Field | D_1 | D_0 | B_1 | B_0 | C_1 | C_0 | A_0 | A_0 | A_1 | A_0 |
| X_1 (Bottom) | X-Frame | A_0 | A_1 | X_0 | N/A | 0 | 0 | A_1 | $A_0 \& A_1$ | A_1 | $A_0 \& A_1$ |
| | X-Field | D_1 | D_1 | B_1 | B_1 | C_1 | C_1 | A_0 | A_1 | A_1 | A_1 |

The final availability tables for each neighbor is given as follows:

| Top MB Availability - B | | | | | | | |
|-------------------------|-----------------|---------------|-----------------------|----------------------|------------------|------------------|------------------------|
| inputs | | | | | | | output |
| Mbaff | current_mb_y[0] | mb_field_flag | current_mbu_top_avail | mb_nb_field_flags[1] | B0 mb_avail_i[4] | B1 mb_avail_i[5] | B mb_nb_avail_intra[2] |
| 0 | na | don't care | na | don't care | 0 | na | 0 |
| 0 | na | don't care | na | don't care | 1 | na | 1 |
| 1 | 0 | 0 | na | don't care | don't care | 0 | 0 |
| 1 | 0 | 0 | na | don't care | don't care | 1 | 1 |
| 1 | 0 | 1 | na | 0 | don't care | 0 | 0 |
| 1 | 0 | 1 | na | 0 | don't care | 1 | 1 |
| 1 | 0 | 1 | na | 1 | 0 | don't care | 0 |
| 1 | 0 | 1 | na | 1 | 1 | don't care | 1 |
| 1 | 1 | 0 | 0 | don't care | don't care | don't care | 0 |
| 1 | 1 | 0 | 1 | don't care | don't care | don't care | 1 |
| 1 | 1 | 1 | don't care | don't care | don't care | 0 | 0 |
| 1 | 1 | 1 | don't care | don't care | don't care | 1 | 1 |



| Top Right MB Availability - C | | | | | | |
|-------------------------------|-----------------|---------------|----------------------|------------------|------------------|------------------------|
| Mbaff | inputs | | | | | output |
| | current_mb_y[0] | mb_field_flag | mb_nb_field_flags[2] | C0 mb_avail_i[2] | C1 mb_avail_i[2] | C Mb_nb_avail_intra[1] |
| 0 | na | don't care | don't care | 0 | na | 0 |
| 0 | na | don't care | don't care | 1 | na | 1 |
| | | | | | | |
| 1 | 0 | 0 | don't care | don't care | 0 | 0 |
| 1 | 0 | 0 | don't care | don't care | 1 | 1 |
| 1 | 0 | 1 | 0 | don't care | 0 | 0 |
| 1 | 0 | 1 | 0 | don't care | 1 | 1 |
| 1 | 0 | 1 | 1 | 0 | don't care | 0 |
| 1 | 0 | 1 | 1 | 1 | don't care | 1 |
| 1 | 1 | 0 | don't care | don't care | don't care | 0 |
| 1 | 1 | 1 | Don't care | don't care | 0 | 0 |
| 1 | 1 | 1 | Don't care | don't care | 1 | 1 |

| Top Left MB Availability - D | | | | | | | | | |
|------------------------------|-----------------|---------------|----------------------|------------------|------------------|----------------------|------------------|------------------|----------------------|
| Mbaff | inputs | | | | | | | | output |
| | current_mb_y[0] | mb_field_flag | mb_nb_field_flags[0] | A0 mb_avail_i[6] | A1 mb_avail_i[7] | mb_nb_field_flags[3] | D0 mb_avail_i[0] | D1 mb_avail_i[1] | Mb_nb_avail_intra[0] |
| 0 | Na | don't care | don't care | don't care | na | don't care | 0 | na | 0 |
| 0 | Na | don't care | don't care | don't care | na | don't care | 1 | na | 1 |
| | | | | | | | | | |
| 1 | 0 | 0 | don't care | don't care | 0 | don't care | don't care | 0 | 0 |
| 1 | 0 | 0 | don't care | don't care | 1 | don't care | don't care | 1 | 1 |
| 1 | 0 | 1 | don't care | don't care | 0 | 0 | don't care | 0 | 0 |
| 1 | 0 | 1 | don't care | don't care | 1 | 0 | don't care | 1 | 1 |
| 1 | 0 | 1 | don't care | don't care | don't care | 1 | 0 | don't care | 0 |
| 1 | 0 | 1 | don't care | don't care | don't care | 1 | 1 | don't care | 1 |
| 1 | 1 | 0 | 0 | 0 | don't care | don't care | don't care | don't care | 0 |
| 1 | 1 | 0 | 0 | 1 | don't care | don't care | don't care | don't care | 1 |
| 1 | 1 | 0 | 1 | 0 | 0 | don't care | don't care | don't care | 0 |
| 1 | 1 | 0 | 1 | 0 | 1 | don't care | don't care | don't care | 1 |
| 1 | 1 | 0 | 1 | 1 | 0 | don't care | don't care | don't care | 0 |
| 1 | 1 | 0 | 1 | 1 | 1 | don't care | don't care | don't care | 1 |
| 1 | 1 | 1 | don't care | don't care | 0 | don't care | don't care | 0 | 0 |
| 1 | 1 | 1 | don't care | don't care | 1 | don't care | don't care | 1 | 1 |



| Left MB (Top Half) Availability - A | | | | | | |
|-------------------------------------|-----------------|---------------|----------------------|------------------|------------------|------------------------|
| inputs | | | | | | output |
| Mbaff | current_mb_y[0] | mb_field_flag | mb_nb_field_flags[0] | A0_mb_avail_i[6] | A1_mb_avail_i[7] | A_Mb_nb_avail_intra[3] |
| 0 | na | don't care | don't care | 0 | na | 0 |
| 0 | na | don't care | don't care | 1 | na | 1 |
| | | | | | | |
| 1 | 0 | 0 | 0 | 0 | don't care | 0 |
| 1 | 0 | 0 | 0 | 1 | don't care | 1 |
| 1 | 0 | 0 | 1 | 0 | 0 | 0 |
| 1 | 0 | 0 | 1 | 0 | 1 | 0 |
| 1 | 0 | 0 | 1 | 1 | 0 | 0 |
| 1 | 0 | 0 | 1 | 1 | 1 | 1 |
| 1 | 0 | 1 | don't care | 0 | don't care | 0 |
| 1 | 0 | 1 | don't care | 1 | don't care | 1 |
| 1 | 1 | 0 | 0 | don't care | 0 | 0 |
| 1 | 1 | 0 | 0 | don't care | 1 | 1 |
| 1 | 1 | 0 | 1 | 0 | 0 | 0 |
| 1 | 1 | 0 | 1 | 0 | 1 | 0 |
| 1 | 1 | 0 | 1 | 1 | 0 | 0 |
| 1 | 1 | 0 | 1 | 1 | 1 | 1 |
| 1 | 1 | 1 | 0 | 0 | don't care | 0 |
| 1 | 1 | 1 | 0 | 1 | don't care | 1 |
| 1 | 1 | 1 | 1 | don't care | 0 | 0 |
| 1 | 1 | 1 | 1 | don't care | 1 | 1 |

| Left MB (Bottom Half) Availability - E | | | | | | |
|--|-----------------|---------------|----------------------|------------------|------------------|------------------------|
| inputs | | | | | | output |
| Mbaff | current_mb_y[0] | mb_field_flag | mb_nb_field_flags[0] | A0_mb_avail_i[6] | A1_mb_avail_i[7] | A_Mb_nb_avail_intra[4] |
| 0 | na | don't care | don't care | 0 | na | 0 |
| 0 | na | don't care | don't care | 1 | na | 1 |
| | | | | | | |
| 1 | 0 | 0 | 0 | 0 | don't care | 0 |
| 1 | 0 | 0 | 0 | 1 | don't care | 1 |
| 1 | 0 | 0 | 1 | 0 | 0 | 0 |
| 1 | 0 | 0 | 1 | 0 | 1 | 0 |
| 1 | 0 | 0 | 1 | 1 | 0 | 0 |
| 1 | 0 | 0 | 1 | 1 | 1 | 1 |
| 1 | 0 | 1 | 0 | don't care | 0 | 0 |
| 1 | 0 | 1 | 0 | don't care | 1 | 1 |
| 1 | 0 | 1 | 1 | 0 | don't care | 0 |
| 1 | 0 | 1 | 1 | 1 | don't care | 1 |
| 1 | 1 | 0 | 0 | don't care | 0 | 0 |
| 1 | 1 | 0 | 0 | don't care | 1 | 1 |
| 1 | 1 | 0 | 1 | 0 | 0 | 0 |
| 1 | 1 | 0 | 1 | 0 | 1 | 0 |
| 1 | 1 | 0 | 1 | 1 | 0 | 0 |
| 1 | 1 | 0 | 1 | 1 | 1 | 1 |
| 1 | 1 | 1 | don't care | don't care | 0 | 0 |
| 1 | 1 | 1 | don't care | don't care | 1 | 1 |



2.8.2.7 MC Kernel Interface Descriptor Table

A set of Kernels are developed to cover all the AVC motion compensated interpolation for reconstructing the display picture. These kernels are loaded into the main memory (by driver) and are referenced by the corresponding interface descriptor. The interface descriptors collectively form an array/table of entries. To decouple the direct referencing of kernels inside the BSD unit, a remapping table is used. The remapping table is indexed by the modified Mb_type. Its content are the pointers to the interface descriptor table.

Two additional flags are derived as follows:

1) MB Subpartition Exist ?

- If at least one 8x8 block within a MB has a subpartition (4x4, 4x8 or 8x4, as indicated by the Sub_mb_type SE), this flag is set to 1; otherwise it is set to 0

2) MB BiPred Exist ?

- If at least one 8x8 block within a MB is marked with BiPred (as indicated by the Sub_Pred_mode SE), this flag is set to 1; otherwise it is set to 0.

| Modified MbType | AVC MbType | | | | | | Remapping Table Entry Index |
|-----------------|---------------------|---------------------|----------------|----------------------|--|--------------------------------------|-----------------------------|
| | I Slice | P Slice | | B Slice | | | |
| | Mb_type SE | Mb_type SE | Sub_mb_type SE | Mb_type SE | MB Subpartition Exist ? (Sub_mb_type SE) | MB BiPred Exist ? (Sub_Pred_mode SE) | |
| intra_16x16 | 1-24 | 6-29 | n/a | 24-47 | n/a | n/a | 0 |
| intra_8x8 | 0 / txformsize8x8=1 | 5 / txformsize8x8=1 | n/a | 23 / txformsize8x8=1 | n/a | n/a | 1 |
| intra_4x4 | 0 / txformsize8x8=0 | 5 / txformsize8x8=0 | n/a | 23 / txformsize8x8=0 | n/a | n/a | 2 |
| intra_pcm | 25 | 30 | n/a | 48 | n/a | n/a | 3 |
| L0_L1_16x16 | n/a | 0, P_skip | n/a | 1, 2 | n/a | n/a | 4 |
| L0_L1_16x8 | n/a | 1 | n/a | 4, 6, 8, 10 | n/a | n/a | 5 |
| L0_L1_8x16 | n/a | 2 | n/a | 5, 7, 9, 11 | n/a | n/a | 6 |
| L1_L0 8x8_nosub | n/a | 3,4 | 0 | 0, 22, B_skip | 0 (0*,inferred*,1,2) | 0 (Direct*,Pred_L0,Pred_L1) | 7 |
| L0_L1_8x8_sub | n/a | 3, 4 | 1, 2, 3 | 0, 22, B_skip | 1 (0*,inferred*,4,5,6,7,10,11) | 0 (Direct*,Pred_L0,Pred_L1) | 8 |



| | | | | | | | |
|--------------|-----|-----|-----|--------------------|--------------------------------|------------------------|----|
| Undefined | n/a | n/a | n/a | n/a | n/a | n/a | 9 |
| BI_16x16 | n/a | n/a | n/a | 3 | n/a | n/a | 10 |
| BI_16x8 | n/a | n/a | n/a | 12, 14, 16, 18, 20 | n/a | n/a | 11 |
| BI_8x16 | n/a | n/a | n/a | 13, 15, 17, 19, 21 | n/a | n/a | 12 |
| BI_8x8_nosub | n/a | n/a | n/a | 0, 22, B_skip | 0 (0*, inferred*, 3) | 1 (Direct*, BiPred) | 13 |
| BI_8x8_sub | n/a | n/a | n/a | 0, 22, B_skip | 1 (0*, inferred*, 8, 9, 12) | 1 (Direct*, BiPred) | 14 |
| Undefined | n/a | n/a | n/a | n/a | n/a | n/a | 15 |

For B MB, in the case of the L1_L0_8x8_nosub (i.e. either L1 or L0 Pred mode for an 8x8 block with no sub-partitioning (i.e. no 4x8, 8x4, or 4x4)), there are 3 possible Mb_type syntax elements (SE), but each may support a different set of sub_mb_type SE and sub_Pred_mode, depending on the direct_8x8_inference_flag and direct mode :

- Mb_type SE = 0 (B_Direct_16x16) and B MB,
 - a. Internal mapped into Mb_type SE = 22 and sub_mb_type[mbPartIdx]=0 (B_Direct_8x8) with
 - i. direct_8x8_inference_flag == 1, spatial direct mode, and only predFlagL0 or predFLagL1 (described in the AVC spec) is set but not both
 - 1. PredL0 or PredL1 8x8 only
- Mb_type SE = 22 (B_8x8) and B MB,
 - a. sub_mb_type[mbPartIdx]=0 (B_Direct_8x8) with
 - i. direct_8x8_inference_flag == 1, spatial direct mode, and only predFlagL0 or predFLagL1 (described in the AVC spec) is set but not both
 - 1. PredL0 or PredL1 8x8 only
 - b. sub_mb_type[mbPartIdx]=1 (B_L0_8x8), =2 (B_L1_8x8)
- Mb_type = inferred B_skip
 - a. Always mapped internally to Mb_type SE = 0 (B_Direct_16x16) , and treated with sub_mb_type[mbPartIdx]=0 (B_Direct_8x8)



For B MB, in the case of the L1_L0_8x8_sub (i.e. either L1 or L0 Pred mode for an 8x8 block with sub-partitioning (i.e. 4x8, 8x4, or 4x4)), there are 3 possible Mb_type syntax elements (SE), but each may support a different set of sub_mb_type SE and sub_Pred_mode, depending on the direct_8x8_inference_flag and direct mode:

- Mb_type SE = 0 (B_Direct_16x16) and B MB,
 - a. Internal mapped into Mb_type SE = 22 and sub_mb_type[mbPartIdx]=0 (B_Direct_8x8) with
 - i. direct_8x8_inference_flag == 0, spatial direct mode, and only predFlagL0 or predFLagL1 (described in the AVC spec) is set but not both
 - 1. PredL0 or PredL1 4x4 only
- Mb_type SE = 22 (B_8x8) and B MB,
 - a. sub_mb_type[mbPartIdx]=0 (B_Direct_8x8) with
 - i. direct_8x8_inference_flag == 0, spatial direct mode, and only predFlagL0 or predFLagL1 (described in the AVC spec) is set but not both
 - 1. PredL0 or PredL1 4x4 only
 - b. sub_mb_type[mbPartIdx]=4,5,6,7,10, and 11 (4x8, 8x4, 4x4 with PredL0 or PredL1)
- Mb_type = inferred B_skip
 - a. Always mapped internally to Mb_type SE = 0 (B_Direct_16x16) , and treated with sub_mb_type[mbPartIdx]=0 (B_Direct_8x8)

For B MB, in the case of the BI_8x8_nosub (i.e. both L1 and L0 Pred modes are in use for an 8x8 block with no sub-partitioning (i.e. no 4x8, 8x4, or 4x4)), there are 3 possible Mb_type syntax elements (SE), but each may support a different set of sub_mb_type SE and sub_Pred_mode, depending on the direct_8x8_inference_flag and direct mode :

- Mb_type SE = 0 (B_Direct_16x16) and B MB,
 - a. Internal mapped into Mb_type SE = 22 and sub_mb_type[mbPartIdx]=0 (B_Direct_8x8) with
 - i. direct_8x8_inference_flag == 1, both direct modes (spatial and temporal) with both predFlagL0 and predFLagL1 (described in the AVC spec) are set
 - 1. BiPred 8x8 only
- Mb_type SE = 22 (B_8x8) and B MB,
 - a. sub_mb_type[mbPartIdx]=0 (B_Direct_8x8) with
 - i. direct_8x8_inference_flag == 1, both direct modes (spatial and temporal) with both predFlagL0 and predFLagL1 (described in the AVC spec) are set
 - 1. BiPred 8x8 only
 - b. sub_mb_type[mbPartIdx]=3 (B_Bi_8x8)



- Mb_type = inferred B_skip
 - a. Always mapped internally to Mb_type SE = 0 (B_Direct_16x16) , and treated with sub_mb_type[mbPartIdx]=0 (B_Direct_8x8)

For B MB, in the case of the BI_8x8_sub (i.e. both L1 or L0 Pred modes are in use for an 8x8 block with sub-partitioning (i.e. 4x8, 8x4, or 4x4)), there are 3 possible Mb_type syntax elements (SE), but each may support a different set of sub_mb_type SE and sub_Pred_mode, depending on the direct_8x8_inference_flag and direct mode :

- Mb_type SE = 0 (B_Direct_16x16) and B MB,
 - a. Internal mapped into Mb_type SE = 22 and sub_mb_type[mbPartIdx]=0 (B_Direct_8x8) with
 - i. direct_8x8_inference_flag == 0, both direct modes (spatial and temporal) with both predFlagL0 and predFlagL1 (described in the AVC spec) are set
 - 1. BiPred 4x4 only
- Mb_type SE = 22 (B_8x8) and B MB,
 - a. sub_mb_type[mbPartIdx]=0 (B_Direct_8x8) with
 - i. direct_8x8_inference_flag == 0, both direct modes (spatial and temporal) with both predFlagL0 and predFlagL1 (described in the AVC spec) are set
 - 1. BiPred 4x4 only
 - b. sub_mb_type[mbPartIdx]=8, 9, 12 (4x8, 8x4, 4x4 with BiPred)
- Mb_type = inferred B_skip
 - a. Always mapped internally to Mb_type SE = 0 (B_Direct_16x16) , and treated with sub_mb_type[mbPartIdx]=0 (B_Direct_8x8)

Internally, after MV and Ref_Idx are generated for the direct modes (B MBs), the remaining processing and MB reconstruction are identical to the regular MB with MVinfo obtained from the bitstream. That is, all “Direct” predmode are internally remapped into PredL0, PredL1 or BiPred.

Although the “Remapping Table Entry Index” ranges from 0 to 15, the number of actual kernels may be lesser. That is, multiple Remapping Indices may refer to the same kernel.



2.8.3 AVC-ILDB Data Buffer Description

2.8.3.1 Deblocking Indices

2.8.3.2 Deblocking Edge Filter Strength

2.8.4 Special Scan order for MB/MB Pair Kernel Processing

MBAFF mode with `special_scan_order = 1`, we will be spawning MBs in this order: (for IT/MC – top of the pairs in the same wave front first and the bottom of the pairs)

| | 0 | 1 | 2 | 3 | 4 |
|----------|----------|----------|----------|----------|----------|
| 0 | 0 | 2 | 4 | 8 | 12 |
| 1 | 1 | 3 | 6 | 10 | 15 |
| 2 | 5 | 9 | 13 | 18 | 22 |
| 3 | 7 | 11 | 16 | 20 | 24 |
| 4 | 14 | 19 | 23 | 26 | 28 |
| 5 | 17 | 21 | 25 | 27 | 29 |

(Performance Based Spawning)

And ILDB, the special scan order, do a pair at a time. The kernel is simpler when the edge control data are layout in this order. The pair can be split in processing if they are not depending on each other.

| | 0 | 1 | 2 | 3 | 4 |
|----------|----------|----------|----------|----------|----------|
| 0 | 0 | 2 | 4 | 8 | 12 |
| 1 | 1 | 3 | 5 | 9 | 13 |
| 2 | 6 | 10 | 14 | 18 | 22 |
| 3 | 7 | 11 | 15 | 19 | 23 |
| 4 | 16 | 20 | 24 | 26 | 28 |
| 5 | 17 | 21 | 25 | 27 | 29 |

(MBs spawned in pairs)



2.9 Error Handling and Error Concealment

Design objective is to keep the playback continue going forward in time even when errors are encountered (detected or not). Hence, the system will reconstruct and display the new picture as much as possible, without dropping the current picture nor repeating the previous decoded picture.

With the limitation of the current architecture and H/W design, error handling is performed only at the Slice boundary, and in between pictures.

To catch the case when the AVC_BSD has finished decoding the current picture but the last MB is not reached, a Phantom Slice is required to conceal such error condition. Therefore, it is required to submit an additional AVC_BSD_OBJECT command for the Phantom Slice at the end of each picture. A Phantom Slice is one that has the indirect_data_length set to 0 and Slice_Start_Mb_Num set to the total number of MB in the current picture + 1.

2.9.1 Design Assumptions

Operating Assumptions:

- The BSD unit is either actively decoding a picture, or in between pictures when errors occur.
- At the beginning of decoding a new picture (the first slice of a picture), the current MB location registers are reset to 0 by the H/W (??? Automatically or through State Command)
- The end of a picture is tracked by comparing the current MB location (just being decoded) to the the picture size parameters (in MB unit) in the BSD unit. The BSD unit will never attempt to decode beyond the picture size programmed, even when there are data remained at the BSD input.
- The end of a picture sequence is only known in the driver. BSD unit will continue operating until no more decoding request commands are injected or when there is an optional time-out by the driver (??? Are we going to have a watchdog timer ???).
- BSD unit has no notion of picture structure, it only deals with Slice as instructed and directed by the State Commands. It does not know it is decoding the last slice of a picture (until starting decoding the next picture ??? is it true).
- There is no H/W Start_Code_Detector at the front end of the BSD unit, so the end of Slice signal can only be detected and generated internally.



2.9.2 Error Concealment and Recovery Strategies

Depending on the category and position of error being detected, a concealment strategy will be adopted.

- Picture Repeat
 - Handled by the display controller – when it does not receive a new picture to display at the appropriate time period
- H/W MB Fillings (as described below)
- S/W Re-sync at the next Slice or picture boundary
 - Dropping a Slice, or
 - The driver can ignore the rest of the bitstream until the start of the next NAL is detected. That is to start decoding again at slice boundary
 - Dropping a picture entirely.
 - The driver can decide to ignore all subsequent slices until a new IDR has arrived, and to start decoding again from there on.

2.9.3 Error Handling

2.9.3.1 S/W Driver

Provide high level error handling capabilities – interrupts and/or polling. Perform error handling and recovery in between Slices and Pictures boundaries.

It also checks for the bitstream syntax errors at and above Slice Header Layers, before proceeding with the Slice decoding.

Currently there is no communication of error handling at the application level.

There are also two error status bits generated in the H/W that the S/W can poll regularly :

- Force_In_Progress
 - A H/W Status Bit (read-only) to indicate filling missing MBs is in progress
- Error_In_Slice
 - A H/W Status Bit (read-only) to indicate an error has detected in the HW, including VLC Decoding Error, BSD Decoded Data Error and MPR Error



2.9.3.2 H/W MB Filling

This is the only error handling mechanism built-in the BSD unit. It will automatically perform error detection and error concealment in the Slice Data Layer at the Slice boundary. There is no direct S/W control.

When decoding a Slice (contains multiple MBs), internally the BSD unit will keep track of the number of MBs being decoded so far. This number also represents the position of the corresponding MB in a picture, assuming a raster-scan MB order (top to bottom, left to right). Hence, the Current_MB_Number register is maintained and updated in the H/W automatically.

For each Slice Data, there is also a Start_MB_Number presented at the corresponding Slice Header. This number specifies the picture location of the very first MB decoded from this Slice Data.

Just before start decoding the current new Slice Data, the last value of the Current_MB_Number register (holding the last MB position of the previously decoded Slice) is compared against the Start_MB_Number. If there is a match, the decoding of current new Slice continues. Otherwise, error concealment is immediately triggered.

There are 2 erroneous cases:

- Current_MB_Number > Start_MB_Number
 - Current_MB_Number will force to equal to Start_MB_Number, and decoding continues at the newly-adjusted Current_MB_Number location.
 - If Current_MB_Number has reached the end of the picture, any data remains in the input must be discarded.
- Current_MB_Number < Start_MB_Number
 - The gap is automatically filled, so that the decoding can continue at the after-filled Current_MB_Number location
 - The starting position for MB filling is adjustable, and is equal to $\text{Current_MB_Number} - \text{MIN}\{\text{Decoded_MB_Count_Prev}, [(MBAFF ? 2:1) * \text{Force_Skip_Rewind}]\}$
 - The MBs used to fill the gap are derived as follows :
 - Intra Concealment Mode
 - If there are no pictures in the DPB, the missing MBs are concealed as if they had been Intra_16x16_2_0_0 and Intra_Chroma_DC MBs
 - Inter Concealment Mode
 - The missing MBs are concealed by copying co-located MBs from the frame store specified as the error concealment frame store during the slice start command. By default the top of the list_0 is used.
 - If the current premature completed slice is the last slice of a picture, SW will need to inject a Phantom Slice to cause the HW to fill the missing MBs to the end of the picture
 - A Phantom Slice with Start_MB_NUMBER set beyond or equal to the picture size
 - Decoding stops when the picture size has reached.



For example: a picture contains 3 Slices

- Slice 1 – contains 100 MBs
- Slice 2 – contains 200 MBs
- Slice 3 – contains 50 MBs

After decoded the Slice Data of Slice 1, the Current_MB_Number is rest at 80, When start to decode Slice 2, its Start_MB_Number is set at 100. Hence, the H/W will automatically fill-in the missing MBs from 81 to 99. Then Slice 2 will continue its decoding from location 100.

Erroneous MB decoding will never pass to the next picture. Since, internally the BSD unit has stored the size of the picture (in MB unit), and it will stop further processing (error fillings, or MB decoding) once this boundary value has reached.

Refer to later section on supporting multiple stream decoding, in which the Current_MB_Number register can be programmed from the driver to allow additional flexibility in error handling.

2.9.3.3 Error Statistics

Another high level error handling mechanism is to keep track of

- the error count in decoding each picture of a video sequence, and/or
- the number of MBs being concealed in a picture

Depending on the allowable thresholds, if the error counts are within tolerances, nothing will be done. Otherwise, actions like – re-establishing the video source linkage, can be taken, or choosing a different error handling processing.

2.9.3.4 WatchDog Timers

At the driver level an additional error detection mechanism is needed. In the situation where errors have stalled the BSD unit from progressing, there must be a way to tell that the H/W is hung. Watch-dog timers are preset with some known timing constraints that are related to the decoding, e.g., an expected decoding time of a slice or a picture. Once the threshold has reached, a stall condition is assumed, and action will take place to reset the BSD unit, and up to the driver to decide where to start decoding again or simply abort the decoding processing.



2.9.4 Error Handling for Non-Existing Reference Picture

There are 3 options:

- For B-MB/Partition Direct Prediction Modes :
 - Spatial Direct Mode
- If co-located MB reference picture does not exist, zero-motion vectors are assumed.
 - Temporal Direct Mode
- If either co-located MB reference picture does not exist, zero-motion vectors are assumed.
- Let the S/W to provide an appropriate replacement/alternative frame store ID
 - Each Reference Picture is addressed by a byte in the MPR List
 - If a list entry reference a picture that does not exist, the value written for that entry should flag it as non-existing
- Bit 7 is the Non-Existing Flag, Bit4:0 will provide the alternate Frame Store Index
- Follow the SEI Recovery Point Conformance (not implement in this version, since we have taken out the Greylevel Non-Existing Reference flag)
 - The recovery point SEI message assists a decoder
 - After decoder initiates random access, or after the encoder indicates a broken link in the sequence
 - Decoder has the option for Inter-Prediction to use mid-level grey for all references to “not available” reference pictures (instead of selecting an alternative frame store). That is references to pictures containing only intra macroblocks and having sample values given by Y equal to 128, Cb equal to 128, and Cr equal to 128 (mid-level grey)

2.10 Concurrent, Multiple Video Stream Decoding Support

The natural place for switching across multiple streams is at the Slice boundary. Each Slice is a self-sustained unit of compressed video data and has no dependency with its neighbors (except for the Deblocking process). In addition, there is no interruptability within a Slice. However, when ILDB is invoked, the processing of some MBs will require neighbour MB information that cross the Slice boundary. Hence, to limit the buffering requirement, in this version of H/W design, the stream switching can only be performed at the picture boundary instead.



2.11 Performance and Latency Estimation

2.12 Media Messages

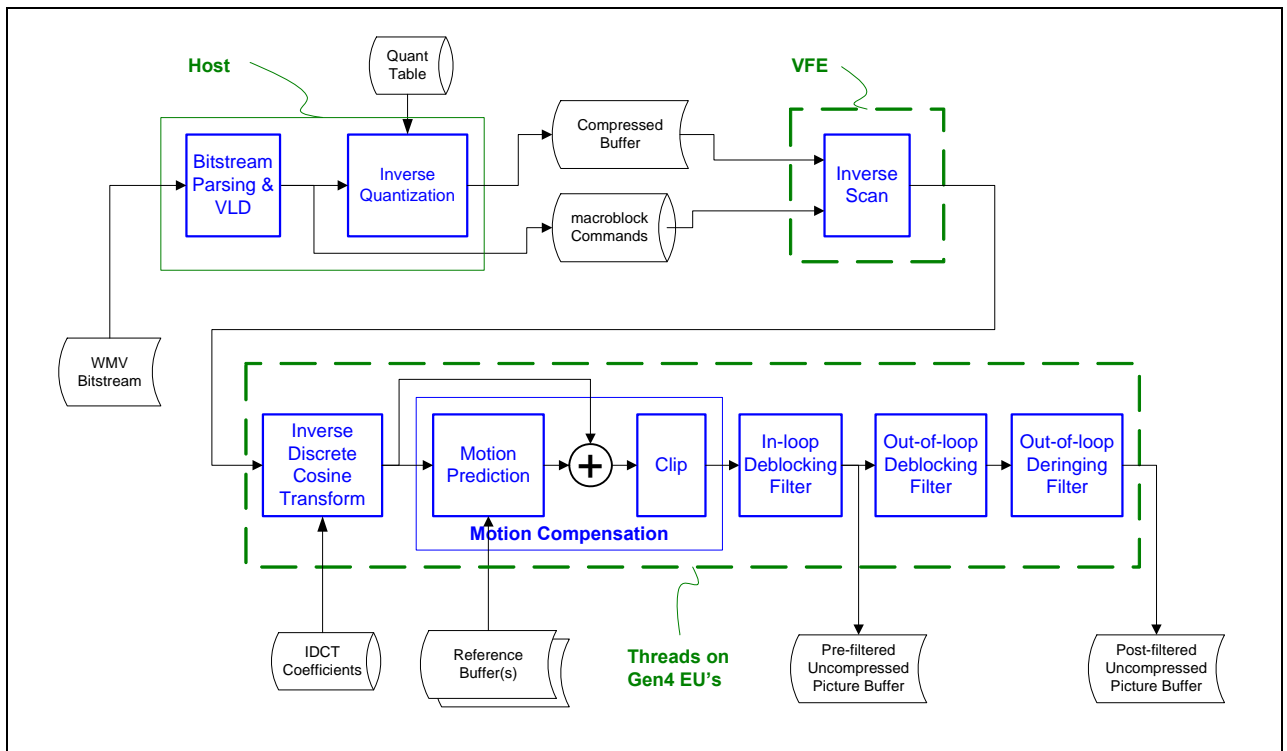
Since this is a standalone bit stream co-processor, no Media Messaging is implemented. It does not have threading capability to the EU Subsystem.

2.13 Media Applications with Specific Hardware Support

2.13.1 AVC – Off-host CABAC/CAVLD Acceleration

Error! Reference source not found. shows the complete AVC video decode pipeline and one of the host/graphics partitions. This partition can be referred to as the off-host CABAC/CAVLD partition as the graphics device is responsible of performance inverse-scan, IDCT, motion compensation and the rest of the post filter operations.

AVC decoding pipeline with off-host CABAC/CAVLD acceleration





3. VC1 Bit-Serial Decoder ([DevCTG, DevILK])

3.1 Introduction

The parallel Media Decode Engine (MDE) also contains a VC1 Bit-Serial Decoding (BSD) module.

- Supporting VC1 Advanced Profile only (no supporting Simple Profile or Main Profile)

3.2 VC1_BSD Commands

3.2.1 VC1_IND_OBJ_BASE_ADDR Command

The VC1_IND_OBJ_BASE_ADDR command sets the memory base address pointers for the subsequent Indirect Data Start Address specified in the VC1_BSD_OBJECT commands.

While the use of this base addresses is unconditional, the indirection can be effectively disabled by setting the base addresses to zero.

The Command Streamer (BCS) will perform the memory access bound check automatically using the VC1 Indirect Object Access Upper Bound specification. If any access is at or beyond this bound, zero value is returned. The request to memory still being sent, but the BSD H/W will detect and perform the zeroing. If the Upper Bound is turned off, the beyond bound request will return whatever on the bus (invalid data).

| Notation | Definition |
|----------------------|---|
| PhysicalAddress[n:m] | Corresponding bits of a physical graphics memory byte address (not mapped by a GTT) |
| GraphicsAddress[n:m] | Corresponding bits of an absolute, virtual graphics memory byte address (mapped by a GTT) |

| Dword | Bit | Description |
|-------|-------|---|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | VC1 Command Opcode = VC1_IND_OBJ_BASE_ADDR Pipeline[28:27] = BSD = 2h; Opcode[26:24] = VC1 = 5h; Sub Opcode[23:16] = 0h |
| | 15:0 | DWord Length (Excludes DWords 0, 1) = 0001h |
| 1 | 31:12 | VC1 Indirect Object Base Address. Specifies the 4K-byte aligned memory base address for indirect object load in VC1_BSD_OBJECT command. Format = GraphicsAddress[31:12] |
| | 11:0 | Reserved : MBZ |



| Dword | Bit | Description |
|-------|-------|---|
| 2 | 31:12 | VC1 Indirect Object Access Upper Bound. This field specifies the 4K-byte aligned (exclusive) maximum Graphics Memory address access by an indirect object load in a VC1_BSD_OBJECT command. Indirect data accessed at this address and beyond will appear to be 0. Setting this field to 0 will cause this range check to be ignored. If non-zero, this address must be greater than the VC1 Indirect Object Base Address . Hardware ignores this field if indirect data is not present. Format = GraphicsAddress[31:12] |
| | 11:0 | Reserved: MBZ |

3.2.2 VC1_BSD_STATE Commands

3.2.2.1 VC1_BSD_PIC_STATE Command

VC1_BSD_PICTURE_STATE command encapsulates the decoding parameters that are read or derived from bitstream syntax elements above (inclusive) picture header layer. These parameters are static for a picture and when slice structure is present, these parameters are not changed from slice to slice of the same picture. Hence, this command is only issued at the beginning of processing a new picture and prior to the VC1_BSD_OBJECT command. The values set for these state variables are retained internally across Slices.

Only the parameters needed by hardware to decode bit sequence for the macroblocks in a picture layer or a slice layer are presented in this command. Other parameters such as the ones used for inverse transform or motion compensation are not here.

| Dword | Bit | Description |
|-------|-------|---|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | VC1 Command Opcode = VC1_BSD_PIC_STATE Pipeline[28:27] = BSD = 2h; Opcode[26:24] = VC1 = 5h; Sub Opcode[23:16] = 1h |
| | 15:0 | DWord Length (Excludes DWords 0,1) = 0005h |
| 1 | 31:24 | Reserved. MBZ. |
| | 23:16 | PictureHeightInMBs (Picture Height in Macroblocks) This field indicates the height of the picture in unit of macroblocks. For example, for a 1920x1080 frame picture, PictureHeightInMBs equals 68 (1080 divided by 16, and rounded up, i.e. effectively specified as 1088 instead). Format: U8 with a valid range of [1, 68] <i>Note: Even though the Advanced Profile allows frame dimensions (width, height) to not be aligned to macroblock boundary, it doesn't affect the bitstream decoding. And it is preferable to use 'intermediate buffer' that is macroblock aligned for decoding. In order to simplify the out-of-bound reference pixel access, the out-of-bound extrapolation rule in VC1 spec can be used to expand the expected decoded frame to the intermediate buffer dimension.</i> |
| | 15:8 | Reserved. MBZ. |



| Dword | Bit | Description |
|-------|-------|---|
| | 7:0 | <p>PictureWidthInMBs (Picture Width in Macroblocks)</p> <p>This field indicates the width of the picture in unit of macroblocks. For example, for a 1920x1080 frame picture, PictureWidthInMBs equals 120 (1920 divided by 16).</p> <p>Format: U8 with a valid range of [1,120]</p> |
| 2 | 31:30 | Reserved. MBZ |
| | 29 | <p>ConcealmentEnable</p> <p>This field when set to 1 indicates HW needs to perform concealment</p> <p>This field should be "1", since error concealment should always occur to fill in any missing MB for the properly functioning of VFE.</p> <p>Other wise no concealment needed.</p> |
| | 28:24 | <p>AltPQuant (Alternative Picture Quantization Value)</p> <p>This field is identical to the variable ALTPQUANT which is derived from VOPDQUANT configuration in the VC1 standard.</p> |
| | 23 | <p>LoopFilter</p> <p>This filed is the decoded syntax element LOOPFILTER in bitstream</p> <p>Indicates if In-loop Deblocking is ON at picture level</p> <p>Disable the output of control information to the VFE pipe when loopfilter is off</p> |
| | 22 | <p>Overlap</p> <p>This filed is the decoded syntax element OVERLAP in bitstream</p> <p>Indicates if Overlap smoothing is ON at picture level</p> |
| | 21 | <p>ImplicitQuantizer</p> <p>This bit is set to 1 when syntax element QUANTIZER=0, else its set to 0</p> |
| | 20:16 | <p>PQuant (Picture Quantization Value)</p> <p>This is the same as the calculated variable PQUANT in VC1 standard where $PQuant = PQINDEX$,</p> <p>except when $QUANTIZER = 0$ and $PQINDEX > 8$, it is given as</p> $PQuant = (PQINDEX < 29) ? PQINDEX - 3 : PQINDEX * 2 - 31$ |
| | 15:12 | <p>AltPQuantEdgeMask (Alternative Picture Quantization Edge Mask)</p> <p>This field is a bit mask for the four edges in clock-wise order, indicating whether AltPQuant is used for the edge macroblocks.</p> <p>It is derived based on the following variables <i>DQUANT</i>, <i>DQUANTFRM</i>, <i>DQPROFILE</i>, <i>DQSBEDGE</i>, <i>DQDBEDGE</i>, and <i>DQBILEVEL</i> defined in the VC1 standard, as shown in Table 3-1.</p> <p>This field is valid only if AltPQuantConfig is 01.</p> <ul style="list-style-type: none"> Bit 0: Left picture edge macroblocks Bit 1: Top picture edge macroblocks Bit 2: Right picture edge macroblocks Bit 3: Bottom picture edge macroblocks |



| Dword | Bit | Description |
|-------|-------|--|
| | 11:10 | <p>AltPQuantConfig (Alternative Picture Quantization Configuration)</p> <p>This field specifies the way AltPQuant is used in the picture. It determines how to compute the macroblock quantizer step size, MQQUANT.</p> <p>It is derived based on the following variables <i>DQUANT</i>, <i>DQUANTFRM</i>, <i>DQPROFILE</i>, <i>DQSBEDGE</i>, <i>DQDBEDGE</i>, and <i>DQBILEVEL</i> defined in the VC1 standard, as shown in Table 3-1.</p> <ul style="list-style-type: none"> 00 = AltPQuant not used 01 = AltPQuant is used and applied to edge macroblocks only 10 = MQQUANT is encoded in macroblock layer 11 = AltPQuant and PQuant are selected on macroblock basis |
| | 9 | <p>HalfQP</p> <p>This field is used for inverse quantization of AC coefficients. It is valid only when PQuant is used.</p> |
| | 8 | <p>PQuantUniform</p> <p>Indicating if uniform quantization applies to the picture. It is used for inverse quantization of the AC coefficients.</p> <ul style="list-style-type: none"> 0 – Non-uniform 1 – Uniform |
| | 7:6 | <p>CondOver</p> <p>This field is the decoded syntax element CONDOVER in bitstream with the following mapping:</p> <ul style="list-style-type: none"> 00 = Overlap transform is not done unless PQUANT > 8 (equivalent to 0b in bitstream/spec) 01 = Reserved (illegal) 10 = Overlap transform is done (equivalent to 10b in bitstream/spec) 11 = Overlap transform is done if OVERFLAGS MB is 1 (equivalent to 11b in bitstream/spec) <p>Note that the mapping used by the reference decoder is different from the above.</p> |
| | 5 | <p>SecondField</p> <p>This field is set for the second field in field pictures.</p> |



| Dword | Bit | Description |
|-------|-------|--|
| | 4:2 | <p>PicType (Picture Type) This field specifies the coding type of the picture according to the Frame Coding Mode.</p> <p>When FCM = 00 01 (a Progressive or Interlaced Frame Picture):</p> <ul style="list-style-type: none"> 000 = I 001 = P 010 = B 011 = BI 100 = Skipped Other encodings are reserved <p>When FCM = 10 11 (a Field Picture)</p> <ul style="list-style-type: none"> 000 = I/I 001 = I/P 010 = P/I 011 = P/P 100 = B/B 101 = B/BI 110 = BI/B 111 = BI/BI |
| | 1:0 | <p>FCM (Frame Coding Mode) This is the same as the variable FCM defined in VC1.</p> <ul style="list-style-type: none"> 00 – Progressive Frame Picture 01 – Interlaced Frame Picture 10 – Field Picture with Top Field First 11 – Field Picture with Bottom Field First. |
| | 31:24 | Reserved. MBZ. |
| 3 | 23 | <p>BitplanePresentFlag (Bitplane Buffer Present Flag) This field indicates whether the bitplane buffer is present for the picture. If set, at least one of the fields listed in bits 22:16 is coded in non-raw mode, and Bitplane Buffer Base Address field in the VC1_BSD_BUF_BASE_STATE command points to the bitplane buffer. Otherwise, all the fields that are applicable for the current picture in bits 22:16 must be coded in raw mode.</p> <ul style="list-style-type: none"> 0 = bitplane buffer is not present 1 = bitplane buffer is present |
| | 22 | <p>ForwardMbRaw This field indicates whether the FORWARDMB field is coded in raw or non-raw mode.</p> <p>This field is only valid when PictureType is B.</p> <ul style="list-style-type: none"> 0 = non-raw mode 1 = raw mode |



| Dword | Bit | Description |
|-------|-------|---|
| | 21 | <p>MvTypeMbRaw</p> <p>This field indicates whether the MVTYPEPREMB field is coded in raw or non-raw mode.</p> <p>This field is only valid when PictureType is P.</p> <p>0 = non-raw mode 1 = raw mode</p> |
| | 20 | <p>SkipMbRaw</p> <p>This field indicates whether the SKIPMB field is coded in raw or non-raw mode.</p> <p>This field is only valid when PictureType is P or B.</p> <p>0 = non-raw mode 1 = raw mode</p> |
| | 19 | <p>DirectMbRaw</p> <p>This field indicates whether the DIRECTMB field is coded in raw or non-raw mode.</p> <p>This field is only valid when PictureType is P or B.</p> <p>0 = non-raw mode 1 = raw mode</p> |
| | 18 | <p>OverflagsRaw</p> <p>This field indicates whether the OVERFLAGS field is coded in raw or non-raw mode.</p> <p>This field is only valid when PictureType is I or BI.</p> <p>0 = non-raw mode 1 = raw mode</p> |
| | 17 | <p>AcPredRaw</p> <p>This field indicates whether the ACPRED field is coded in raw or non-raw mode.</p> <p>This field is only valid when PictureType is I or BI.</p> <p>0 = non-raw mode 1 = raw mode</p> |
| | 16 | <p>FieldTxRaw</p> <p>This field indicates whether the FIELDTX field is coded in raw or non-raw mode.</p> <p>This field is only valid when PictureType is I or BI.</p> <p>0 = non-raw mode 1 = raw mode</p> |
| | 15:14 | <p>ExtendedDMVRange (Extended Differential Motion Vector Range Flag)</p> <p>This field specifies the differential motion vector range in interlaced pictures. It is equivalent to the variable <i>DMVRANGE</i> in the VC1 standard.</p> <p>00 = No extended range 01 = Extended horizontally 10 = Extended vertically 11 = Extended in both directions</p> |



| Dword | Bit | Description |
|-------|-------|---|
| | 13:12 | <p>ExtendedMVRRange (Extended Motion Vector Range Flag)</p> <p>This field specifies the motion vector range in quarter-pel or half-pel modes. It is equivalent to the variable <i>MVRANGE</i> in the VC1 standard.</p> <p>00 = [-256, 255] x [-128, 127] 01 = [-512, 511] x [-256, 255] 10 = [-2048, 2047] x [-1024512, 1023511] 11 = [-4096, 4095] x [-20481024, 20471023]</p> |
| | 11 | <p>FourMvSwitch (Four Motion Vector Switch)</p> <p>This field indicates if 4-MV is present for an interlaced frame P picture. It is identical to the variable <i>4MVSWITCH</i> (<i>4 Motion Vector Switch</i>) in VC1 standard.</p> <p>0 = only 1-MV 1 = 1, 2, or 4 MVs</p> |
| | 10 | <p>FastUVMCFlag (Fast UV Motion Compensation Flag)</p> <p>This field specifies whether the motion vectors for UV is rounded to half or full pel position. It is identical to the variable <i>FASTUVMC</i> in VC1 standard.</p> <p>0 = no rounding 1 = quarter-pel offsets to half/full pel positions</p> |
| | 9:8 | <p>UnifiedMvMode (Unified Motion Vector Mode)</p> <p>This field is a combination of the variables <i>MVMODE</i> and <i>MVMODE2</i> in the VC1 standard.</p> <p>11 = 1-MV half-pel bilinear 01 = 1-MV 10 = 1-MV half-pel 00 = Mixed MV</p> |
| | 7 | Reserved. MBZ. |
| | 6 | <p>RefFieldPicPolarity (Reference Field Picture Polarity)</p> <p>This field specifies the polarity of the one reference field picture used for a field P picture. It is derived from the variable <i>REFFIELD</i> defined in VC1 standard and is only valid when one field is referenced (<i>NUMREF</i> = 0) for a field P picture.</p> <p>When <i>NUMREF</i> = 0 and <i>REFFIELD</i> = 0, this field is the polarity of the reference I/P field that is temporally closest; When <i>NUMREF</i> = 0 and <i>REFFIELD</i> = 1, this field is the polarity of the reference I/P field that is the second most temporally closest. The distance is measured based on display order but ignoring the repeated field if present (due to <i>RFF</i> = 1).</p> <p>0 = Top (even) field 1 = Bottom (odd) field</p> |
| | 5 | <p>NumRef (Number of References)</p> <p>This field indicates how many reference fields are referenced by the current (field) picture. It is identical to the variable <i>NUMREF</i> in the VC1 standard.</p> <p>This field is only valid for field P picture (<i>FCM</i> = 10 11).</p> <p>0 = One field referenced 1 = Two fields referenced</p> |



| Dword | Bit | Description | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|-------|---------------------|---|-------------|---------------------|-----------------|-------------|---|-----|-----|-----|---|-----|-----|----|---|-----|-----|-----|---|-----|-----|----|---|-----|-----|-----|---|-----|-----|----|---|-----|-----|-----|---|---------|-----|-----|---|---------|-----|-----|----|---------|-----|----|----|---------|-----|-----|----|---------|-----|----|----|---------|-----|----|----|---------|-----|-----|----|---------|-----|-----|----|---------|-----|-----|----|---------|-----|-----|----|---------|-----|----|----|---------|-----|----|----|---------|-----|-----|----|---------|-----|-----|
| | 4:0 | <p>FwdRefDist (Reference Distance)</p> <p>This field is the number of frames between the current frame and its reference frame. It is identical to the variable REFDIST (P Reference Distance) in the VC1 standard. 0 means that the previous frame is the reference frame.</p> <p>This field is also valid for B field pictures giving the value of FRFD</p> <p>Format = U5 with a valid range of [0, 16]</p> | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 31:24 | <p>ScaleFactor</p> <p>This field is the scale factor for computing Direct-mode motion vectors. It is derived from the variable <i>BFACTION</i> in the VC1 standard, section 8.4.5.4.</p> <p>There are only 21 valid values corresponding to the 21 encodings of <i>BFACTION</i> as shown in the table here. Other values are reserved.</p> <p>MSB of this field can be used to determine if <i>BFACTION</i> is greater than or equal to 1/2, which is used to determine Motion Prediction Type for B pictures. Effectively, condition "<i>BFACTION</i> >= 1/2" is equivalent to condition "ScaleFactor >= 128".</p> <p>This field is only valid for B pictures.</p> <table border="1"> <thead> <tr> <th>Item</th> <th><i>BFACTION</i> VLC</th> <th><i>BFACTION</i></th> <th>ScaleFactor</th> </tr> </thead> <tbody> <tr><td>1</td><td>000</td><td>1/2</td><td>128</td></tr> <tr><td>2</td><td>001</td><td>1/3</td><td>85</td></tr> <tr><td>3</td><td>010</td><td>2/3</td><td>170</td></tr> <tr><td>4</td><td>011</td><td>1/4</td><td>64</td></tr> <tr><td>5</td><td>100</td><td>3/4</td><td>192</td></tr> <tr><td>6</td><td>101</td><td>1/5</td><td>51</td></tr> <tr><td>7</td><td>110</td><td>2/5</td><td>102</td></tr> <tr><td>8</td><td>1110000</td><td>3/5</td><td>153</td></tr> <tr><td>9</td><td>1110001</td><td>4/5</td><td>204</td></tr> <tr><td>10</td><td>1110010</td><td>1/6</td><td>43</td></tr> <tr><td>11</td><td>1110011</td><td>5/6</td><td>215</td></tr> <tr><td>12</td><td>1110100</td><td>1/7</td><td>37</td></tr> <tr><td>13</td><td>1110101</td><td>2/7</td><td>74</td></tr> <tr><td>14</td><td>1110110</td><td>3/7</td><td>111</td></tr> <tr><td>15</td><td>1110111</td><td>4/7</td><td>148</td></tr> <tr><td>16</td><td>1111000</td><td>5/7</td><td>185</td></tr> <tr><td>17</td><td>1111001</td><td>6/7</td><td>222</td></tr> <tr><td>18</td><td>1111010</td><td>1/8</td><td>32</td></tr> <tr><td>19</td><td>1111011</td><td>3/8</td><td>96</td></tr> <tr><td>20</td><td>1111100</td><td>5/8</td><td>160</td></tr> <tr><td>21</td><td>1111101</td><td>7/8</td><td>224</td></tr> </tbody> </table> | Item | <i>BFACTION</i> VLC | <i>BFACTION</i> | ScaleFactor | 1 | 000 | 1/2 | 128 | 2 | 001 | 1/3 | 85 | 3 | 010 | 2/3 | 170 | 4 | 011 | 1/4 | 64 | 5 | 100 | 3/4 | 192 | 6 | 101 | 1/5 | 51 | 7 | 110 | 2/5 | 102 | 8 | 1110000 | 3/5 | 153 | 9 | 1110001 | 4/5 | 204 | 10 | 1110010 | 1/6 | 43 | 11 | 1110011 | 5/6 | 215 | 12 | 1110100 | 1/7 | 37 | 13 | 1110101 | 2/7 | 74 | 14 | 1110110 | 3/7 | 111 | 15 | 1110111 | 4/7 | 148 | 16 | 1111000 | 5/7 | 185 | 17 | 1111001 | 6/7 | 222 | 18 | 1111010 | 1/8 | 32 | 19 | 1111011 | 3/8 | 96 | 20 | 1111100 | 5/8 | 160 | 21 | 1111101 | 7/8 | 224 |
| Item | <i>BFACTION</i> VLC | <i>BFACTION</i> | ScaleFactor | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 1 | 000 | 1/2 | 128 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2 | 001 | 1/3 | 85 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 3 | 010 | 2/3 | 170 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4 | 011 | 1/4 | 64 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 5 | 100 | 3/4 | 192 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 6 | 101 | 1/5 | 51 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 7 | 110 | 2/5 | 102 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 8 | 1110000 | 3/5 | 153 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 9 | 1110001 | 4/5 | 204 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 10 | 1110010 | 1/6 | 43 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 11 | 1110011 | 5/6 | 215 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 12 | 1110100 | 1/7 | 37 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 13 | 1110101 | 2/7 | 74 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 14 | 1110110 | 3/7 | 111 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 15 | 1110111 | 4/7 | 148 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 16 | 1111000 | 5/7 | 185 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 17 | 1111001 | 6/7 | 222 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 18 | 1111010 | 1/8 | 32 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 19 | 1111011 | 3/8 | 96 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 20 | 1111100 | 5/8 | 160 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 21 | 1111101 | 7/8 | 224 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |



| Dword | Bit | Description |
|-------|-------|--|
| 4 | 23 | Reserved. MBZ. |
| | 22:20 | <p>MvTab (Motion Vector Table)</p> <p>This field specifies which motion vector table(s) is (are) used for motion vector (differential) decoding in a P or B picture. This field is the combination of the variables <i>MVTAB</i> and <i>IMVTAB</i> in the VC1 standard. Two bits are defined for progressive frame pictures; And two or three bits are defined for interlaced field/frame pictures depending on NUMREF and P or B picture types.</p> <p>This field is valid for P and B pictures. It is not valid for I pictures.</p> <p>For P or B progressive frame pictures</p> <ul style="list-style-type: none"> 0 = Motion Vector Differential VLD Table 0 1 = Motion Vector Differential VLD Table 1 2 = Motion Vector Differential VLD Table 2 3 = Motion Vector Differential VLD Table 3 <p>The other encodings are reserved</p> <p>For P interlace field pictures with NUMREF = 0 or P/B interlace frame pictures</p> <ul style="list-style-type: none"> 0 = 1-Reference Table 0 1 = 1-Reference Table 1 2 = 1-Reference Table 2 3 = 1-Reference Table 3 <p>The other encodings are reserved</p> <p>For P interlace field picture with NUMREF = 1 or B interlaced field pictures</p> <ul style="list-style-type: none"> 0 = 2-Reference Table 0 1 = 2-Reference Table 1 2 = 2-Reference Table 2 3 = 2-Reference Table 3 4 = 2-Reference Table 4 5 = 2-Reference Table 5 6 = 2-Reference Table 6 7 = 2-Reference Table 7 <p>The other encodings are reserved</p> |
| | 19:18 | <p>FourMvBpTab (4-MV Block Pattern Table)</p> <p>This field specifies which table is used to decode the 4-MV block pattern (4MVBP) syntax element in 4-MV macroblocks. It is identical to the variables <i>4MVBPTAB</i> in the VC1 standard, section 9.1.1.37.</p> <p>This field is valid only in interlace frame P, B pictures, or interlace field P, B pictures. It is not valid for I picture.</p> <p>For interlace field P and B pictures, it is only valid if UnifiedMvMode is equal to Mixed-MV Type.</p> <p>For interlace frame P picture, it is only valid if FourMvSwitch is 1.</p> <p>For interlace frame B picture, it is always valid.</p> <ul style="list-style-type: none"> 0 = 4MVBP Table 0 1 = 4MVBP Table 1 2 = 4MVBP Table 2 3 = 4MVBP Table 3 |



| Dword | Bit | Description |
|-------|-------|--|
| | 17:16 | <p>TwoMvBpTab (2MV Block Pattern Table)</p> <p>This field specifies which table is used to decode the 2MV block pattern (2MVBP) syntax element in 2MV field macroblocks. It is identical to the variables <i>2MVBPTAB</i> in the VC1 standard, section 9.1.1.36.</p> <p>This field is valid only in interlace frame P/B pictures. It is not valid for I picture, nor for interlace field P or B pictures.</p> <p>0 = 2MVBP Table 0 1 = 2MVBP Table 1 2 = 2MVBP Table 2 3 = 2MVBP Table 3</p> |
| | 15:14 | Reserved : MBZ |
| | 13:12 | <p>TransType (Picture-level Transform Type)</p> <p>This field specifies the Transform Type at picture level. It is identical to the variable <i>TTFRM</i> in the VC1 standard, section 7.1.1.41.</p> <p>This field is only valid when TransTypeMbFlag is 1. Otherwise, it is reserved and MBZ.</p> <p>This field is set to 00 when <i>VSTRANSFORM</i> is 0 in the entry point layer.</p> <p>00 = 8x8 Transform 01 = 8x4 Transform 10 = 4x8 Transform 11 = 4x4 Transform</p> |
| | 11 | <p>TransTypeMbFlag (Macroblock Transform Type Flag)</p> <p>This field indicates whether Transform Type is fixed at picture level or variable at macroblock level. It is identical to the variable <i>TTMBF</i> in the VC1 standard, section 7.1.1.40.</p> <p>This field is set to 1 when <i>VSTRANSFORM</i> is 0 in the entry point layer.</p> <p>0 = variable transform type in macroblock layer 1 = use picture level transform type TransType</p> |



| Dword | Bit | Description |
|-------|------|---|
| | 10:8 | <p>MbModeTab (Macroblock Mode Table)</p> <p>This field signals which code table is used to decode the macroblock mode syntax element (MBMODE) in the macroblock layer in a P or B picture. This field is identical to the variables <i>MBMODETAB</i> in the VC1 standard, section 9.1.1.33.</p> <p>This field is valid for interlace frame P, B picture and interlace field P, B picture. It is not valid for I picture, nor progressive frame P, B pictures.</p> <p>Two bits are defined for interlace frame P, B pictures; And three bits are defined for interlaced field P, B pictures.</p> <p>Two bits are defined for interlace frame P, B pictures. There are two set of code tables selected based on if UnifiedMvMode is equal to 4-MV Type or not.</p> <ul style="list-style-type: none"> 0 = Code Table 0 1 = Code Table 1 2 = Code Table 2 3 = Code Table 3 <p>Other encodings are invalid</p> <p>Three bits are defined for interlace field P, B pictures. There are two set of code tables selected based on if UnifiedMvMode is equal to Mixed-MV Type or not.</p> <ul style="list-style-type: none"> 0 = Code Table 0 1 = Code Table 1 2 = Code Table 2 3 = Code Table 3 4 = Code Table 4 5 = Code Table 5 6 = Code Table 6 7 = Code Table 7 |
| | 7:6 | <p>TransAcY (Picture-level Transform Luma AC Coding Set Index)</p> <p>This field, together with PQINDEX, specifies which intra AC coding set to be used for decoding the non-zero AC coefficients in a coded luma (Y) block. This field is the combination of the variables <i>TRANSACFRM</i> and <i>TRANSACFRM2</i> in the VC1 standard.</p> <p>For I pictures, TransAcY is the same as <i>TRANSACFRM2</i>. For other pictures, it is the same as <i>TRANSACFRM</i>, and therefore must be programmed to be the same as TransAcUV.</p> <p>This field is valid for all picture types.</p> <ul style="list-style-type: none"> 0 = Coding set index 0 1 = Coding set index 1 2 = Coding set index 2 3 is invalid |



| Dword | Bit | Description |
|-------|-------------|---|
| | 5:4 | <p>TransAcUV (Picture-level Transform Chroma AC Coding Set Index)</p> <p>This field, together with PQINDEX, specifies which inter AC coding sets to be used for decoding the non-zero AC coefficients in a coded chroma (Cb and Cr) block. This field is identical to the variable <i>TRANSACFRM</i> in the VC1 standard.</p> <p>This field is valid for all picture types.</p> <p>0 = Coding set index 0 1 = Coding set index 1 2 = Coding set index 2 3 is invalid</p> |
| | 3 | <p>TransDcTab (Intra Transform DC Table)</p> <p>This field specifies whether the low motion tables or the high motion tables are used to decode the Transform DC coefficients in intra-coded blocks. This field is identical to the variable <i>TRANSDCTAB</i> in the VC1 standard, section 8.1.1.2.</p> <p>This field is valid for all picture types.</p> <p>0 = The high low motion tables 1 = The low high motion tables</p> |
| | 2:0 | <p>CbpTab (Coded Block Pattern Table)</p> <p>This field specifies the table used to decode the CBPCY syntax element for each coded macroblock in P and B pictures. This field is combination of the variable <i>CBPTAB</i> for P and B frame pictures and the variable <i>ICBPTAB</i> in interlace field P, B pictures and interlace frame P, B pictures in the VC1 standard (Table 52 and Table 102).</p> <p>This field is reserved and MBZ for I or BI pictures as I only has a fixed table.</p> <p>000 = Table 0 (Table 169 for P, B frames or Table 124 otherwise) 001 = Table 1 (Table 170 for P, B frames or Table 125 otherwise) 010 = Table 2 (Table 171 for P, B frames or Table 126 otherwise) 011 = Table 3 (Table 172 for P, B frames or Table 127 otherwise) 100 = Table 4 (Table 128 for interlace field/frame P, B pictures) 101 = Table 5 (Table 129 for interlace field/frame P, B pictures) 110 = Table 6 (Table 130 for interlace field/frame P, B pictures) 111 = Table 7 (Table 131 for interlace field/frame P, B pictures)</p> |
| | 31:164 8 | <p>Reserved : MBZ</p> |
| 5 | 15 | <p>Thread Synchronization Overwrite. This field indicates whether hardware overwrites the Thread Synchronization field in the MEDIA_OBJECT_EX command of each output macroblock. When this field is set, the Thread Synchronization field is set only for intra macroblocks (and at least one of the neighbors top, top-left, top-right and left MB is available). Otherwise (if a macroblock is a P, B), the field is not set.</p> <p>0 = Not Overwrite 1 = Overwrite</p> <p>It is used to generate the thread control field in the BSD output to VFE.</p> |



| Dword | Bit | Description |
|-------|-------|--|
| | 14 | <p>Thread Synchronization. This field indicates whether the thread is a Synchronized Root Thread (SRT). When Thread Synchronization Overwrite is not set, this field is forwarded to the corresponding field in the MEDIA_OBJECT_EX command of each output macroblock. When Thread Synchronization Overwrite is set, hardware ignores this field.</p> <p>0 = Not a SRT 1 = SRT</p> <p>It is used to generate the thread control field in the BSD output to VFE.</p> |
| | 13 | <p>DmvSurfaceValid</p> <p>Indicated when the DMV read surface is valid. This surface stored the direct motion vectors.</p> <p>This field is set for B pictures that can refer to a previous P picture for DMV. If there is a I picture before a B (in decoding order) then this field is not set.</p> |
| | 12:8 | <p>BwdRefDist (Reference Distance)</p> <p>This field is valid only in B field pictures giving the value of BRFD.</p> <p>Format = U5 with a valid range of [0, 16]</p> |
| | 7 | <p>Reserved: MBZ</p> |
| 6 | 6:0 | <p>Bitplane Buffer Pitch: Specifies the bitplane buffer pitch in (#Bytes - 1). The pitch can be any multiple of bytes, but must be greater than or equal to PictureWidthInMBs/2.</p> <p>This field should be specified equal to PictureWidthInMBs/2 for better performance</p> <p>Format = U7 pitch in (Bytes - 1).</p> |
| | 31:16 | <p>VC1-IT Command Header, Vc1ItCmdHeader</p> <p>It allows the driver to directly specify the DW0 of the VC1-IT Command to be generated by the BSD unit that will pass down to drive the VFE unit.</p> <p>This parameter is specified for Intel interface only.</p> <p>cmdType[31:29]&MediaCmdOp[28:16]&Dw</p> |
| | 15:0 | <p>MBZ (BSD HW will pass DW length in this field for VFE pipe)</p> |
| | 31:0 | <p>MBZ---VFE Control in VC1-IT mode</p> |
| 7 | | |

AltPQuantConfig and **AltPQuantEdgeMask** are derived based on the following variables *DQUANT*, *DQUANTFRM*, *DQPROFILE*, *DQSBEDGE*, *DQDBEDGE*, and *DQBILEVEL* defined in the VC1 standard, as shown in the following table.



Table 3-1 Definition of AltPQuantConfig and AltPQuantEdgeMask

| Inputs | | | | | | Outputs | | Description |
|--------|-----------|-----------|----------|----------|-----------|-----------------|-------------------|---|
| DQUANT | DQUANTFRM | DQPROFILE | DQDBEDGE | DQSBEDGE | DQBILEVEL | AltPQuantConfig | AltPQuantEdgeMask | |
| 0 | - | - | - | - | - | 00b | 0000b | No AltPQuant |
| 1 | 0 | - | - | - | - | 00b | 0000b | No AltPQuant |
| 1 | 1 | 11b | - | - | 0 | 10b | 0000b | All MBs are different with <i>MQDIFF</i> and <i>ABSMQ</i> |
| 1 | 1 | 11b | - | - | 1 | 11b | 0000b | All MBs may switch with 1-bit <i>MQDIFF</i> |
| 2 | - | - | - | - | - | 01b | 1111b | All edge MBs |
| 1 | 1 | 00b | - | - | - | 01b | 1111b | All edge MBs |
| 1 | 1 | 01b | 00b | - | - | 01b | 0011b | Left and top MBs |
| 1 | 1 | 01b | 01b | - | - | 01b | 0110b | Top and right MBs |
| 1 | 1 | 01b | 10b | - | - | 01b | 1100b | Right and bottom MBs |
| 1 | 1 | 01b | 11b | - | - | 01b | 1001b | Bottom and left MBs |
| 1 | 1 | 10b | - | 00b | - | 01b | 0001b | Left MBs |
| 1 | 1 | 10b | - | 01b | - | 01b | 0010b | Top MBs |
| 1 | 1 | 10b | - | 10b | - | 01b | 0100b | Right MBs |
| 1 | 1 | 10b | - | 11b | - | 01b | 1000b | Bottom MBs |



3.2.2.2 VC1_BSD_BUF_BASE_STATE Command

These are the base addresses for memory buffers allocated for VC1 BSD operations. It does not include the byte stream read address, which is to be specified in the VC1_BSD_OBJECT Command.

When switching video stream, these base address must be reprogrammed.

| Dword | Bit | Description |
|-------|-------|---|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | VC1 Command Opcode = VC1_BSD_BUF_BASE_STATE Pipeline[28:27] = BSD = 2h; Opcode[26:24] = VC1 = 5h; Sub Opcode[23:16] = 3h |
| | 15:0 | DWord Length (Excludes DWords 0,1) = 0044h |
| 1 | 31:0 | Bitplane Read Buffer Base Address This field must be dword aligned. |
| 2 | 31:6 | Row Store Read/Write (Scratch) Buffer Base Address This field must be 64-bytes aligned |
| | 5:0 | Reserved : MBZ |
| 3 | 31:6 | DMV Read Buffer Base Address (Combined DMV0 and DMV1 into one DMV) This field must be 64-bytes aligned This surface is valid for B pictures (in other document it is said both I and B pictures for error concealment in I ???) |
| | 5:0 | Reserved : MBZ |
| 4 | 31:6 | VC1 IT Command Write Buffer Base Address This field must be 64-bytes aligned |
| | 5:0 | Reserved : MBZ |
| 5 | 31:6 | VC1 IT Data Write Buffer Base Address This is the 4KB aligned portion of the base address for the VC-1-IT data write buffer. Together with the VC-1-IT Data Write Offset, it forms the base pointer where the VC-1-IT data (including for example motion vectors, non-zero residual data), where data for the first macroblock will be written to. This field affect the final memory address computed for the VC-1-IT Data output. However, it doesn't affect the Indirect Data Start Address computed by BSD hardware that is in the VC-1-IT Command output (DW3 of the output VC-1-IT Command). Each macroblock within the VC-1-IT data buffer has a fixed size even though only valid data are written to memory. The remaining bytes for each macroblock are undefined. Format = GraphicsAddress[31:12] <i>Programming Note: As the VC-1-IT data buffer will be used by GPE as the indirect object buffer for MEDIA_OBJECT_EX command, this field should be programmed the same as the Indirect Object Base Address field of the STATE_BASE_ADDRESS command.</i> |
| | 5:0 | Reserved : MBZ |



| Dword | Bit | Description |
|-------|------|---|
| 6 | 31:6 | VC1 ILDB Data Write Buffer Base Address This field must be 64-bytes aligned |
| | 5:0 | Reserved : MBZ |
| 7 | 31:6 | DMV Write Buffer Base Address This field must be 64-bytes aligned This surface is valid for P pictures |
| | 5:0 | Reserved : MBZ |

3.2.2.2.1 Bitplane Buffer

Bitplane coding is used in seven different cases in VC-1, although not all the seven syntax elements are present in the same picture header at the same time. The following list shows which syntax elements are coded as bitplanes in each picture header:

- Progressive I and BI picture headers in AP: ACPRED, OVERFLAGS
Field interlace I and BI picture headers in AP: ACPRED, OVERFLAGS
Frame interlace I and BI picture headers in AP: FIELDTX, ACPRED, OVERFLAGS
- Frame interlace P picture headers in AP: SKIPMB
Progressive P picture headers in SP and MP: MVTYPEMB, SKIPMB
Progressive P picture headers in AP: MVTYPEMB, SKIPMB
- Field interlace B picture headers in AP: FORWARDMB
Frame interlace B picture headers in AP: DIRECTMB, SKIPMB
Progressive B picture headers in AP: DIRECTMB, SKIPMB
Progressive B picture headers in MP: DIRECTMB, SKIPMB

There are also seven different modes of coding the bitplane information. Except when the bitplane is coded in raw mode, the bitplane is decoded by the host and provided to the hardware in the bitplane buffer.

Since at most three bitplanes are encoded in any picture header, instead of using a complete byte for signaling the values of all the seven possible bitplanes for each MB, a more efficient approach is used with each byte divided in two nibbles and each nibble carries the data of up to four bitplanes for one MB.

| PictureType | Bits 3, 7 | Bit 2, 6 | Bits 1, 5 | Bits 0, 4 |
|----------------|-----------|-----------|-----------|-----------|
| I or BI | 0 | OVERFLAGS | ACPRED | FIELDTX |
| P | 0 | MVTYPEMB | SKIPMB | 0 |
| B | 0 | FORWARDMB | SKIPMB | DIRECTMB |



The bytes containing the above defined nibbles are stored in the bitplane buffer in raster scan order. The bitplane buffer is a linear buffer with a buffer pitch (as defined by Bitplane Buffer Pitch field in VC1_BSD_PIC_STATE command). When the number of macroblock in a row is odd, the last byte of the row containing the last macroblock in bits 0-3. The first macroblock of the next row starts at the next pitch offset from the first macroblock of the current row.

The bitplane buffer structure must be sent once per picture only if there is one or more syntax elements coded as bitplanes in the picture header.

3.2.2.2.2 VC1 IT Command Buffer

See Media Chapter for “Inline Data Format in VC1-IT Mode”.

3.2.2.2.3 VC1 IT Residual Data Buffer

See Media Chapter for “Indirect Data Format in VC1-IT Mode”.

3.2.2.2.4 VC1 ILDB Data Buffer

3.2.3 VC1_BSD_OBJECT Command(Pre-DevILK)

The VC1_BSD_OBJECT command is the only primitive command for the VC1_BSD Unit. The macroblock data portion of the bitstream is loaded as indirect data object.

Before issuing an AVC_BSD_OBJECT command, all BSD states need to be valid. Therefore the commands used to set these states need to have been issued at some point prior to the issue of an AVC_BSD_OBJECT command.

VC1 deblock filter kernel cross the slice boundary if in the last MB row of a slice, so need to know the last MB row of a slice to disable the edge mask. There is why VC1 BSD H/W need to know the end of MB address for the current slice. As such no more phantom slice is needed for VC1, as long as the driver will program both start MB address in the current slice and the start MB address of the next slice. As a result, we can also support multiple picture state commands in between slices.

To handle encrypted bitstream decoding, the driver is required to align the indirect bitstream data to a 16byte address boundary, the first byte of the encrypted bitstream data. The application should ensure that after the last valid byte of bitstream data and before encryption, enough padding is added up to the end of a 16byte chunk. In addition, for a protected context, all kernel source and destination data surface addresses must be in PCM space.

| Dword | Bits | Description |
|-------|-------|--|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | VC1 Command Opcode = VC1_BSD_OBJECT Pipeline[28:27] = BSD = 2h; Opcode[26:24] = VC1 = 5h; Sub Opcode[23:16] = 8h |
| | 15:0 | DWord Length (Excludes DWords 0,1) = 0004h |



| Dword | Bits | Description |
|-------|-------|---|
| 1 | 31 | Bit stream is AES-128 Counter mode Encrypted '0' – Not encrypted '1' - Encrypted |
| | 30:22 | Reserved. MBZ |
| | 21:0 | Indirect Data Length. This field provides the length in bytes of the indirect data. A value zero indicates that indirect data fetching is disabled – subsequently, the Indirect Data Start Address field is ignored. This field must have the same alignment as the Indirect Object Data Start Address. 22 bits support the maximum bitstream buffer size of (e.g. for 1080i) is $3072 \times 120 \times 68 / 8 = 3133440$ bytes, which meets the requirements of the VC1 Specification. It includes the byte that contains the First MB Bit Offset. Format = U22 in bytes |
| 2 | 31:29 | Reserved : MBZ |
| | 28:0 | Indirect Data Start Address. This field specifies the Graphics Memory starting address of the data to be fetched into BSD Unit for processing. This pointer is relative to the MDE Indirect Object Base Address . It is a byte-aligned address for the VC1 bitstream data. Range = [0 - 512MB] |
| 3 | 31:30 | Reserved : MBZ |
| | 31:24 | SliceStartVertPos (Slice Vertical Position) This field specifies the position in y-direction of the first macroblock in the Slice in unit of macroblocks. For SecondField this value is reset to zero as opposed to the VC1 spec Ref: 9.1.2 Slice Layer. |
| | 23:3 | Reserved : MBZ |
| | 2:0 | FirstMbBitOffset (First Macroblock Bit Offset) This field provides the bit offset of the first macroblock of the Slice in the first byte of the input compressed bitstream. Format: U3 |
| 4 | 31:30 | Reserved : MBZ |
| | 31:24 | NextSliceStartVertPos (Next Slice Vertical Position) This field specifies the position in y-direction of the first macroblock in the next Slice in unit of macroblocks. This field is primarily used for error concealment. In the case that current slice is the last slice, this field should set to the height of picture (since y-direction is zero-based numbering) |
| | 23:0 | Reserved : MBZ |
| 5 | 31:0 | Initial counter value for AES-128 Counter mode decryption 32-bit initial counter value (which does not need to be secure) for the AES counter mode decryption. If Bit stream is AES-128 Counter mode Encrypted is 0 , i.e. not encrypted bit stream, this field is ignored and should be MBZ. |



3.2.4 VC1_BSD_OBJECT Command(DevILK only)

The VC1_BSD_OBJECT command is the only primitive command for the VC1_BSD Unit. The macroblock data portion of the bitstream is loaded as indirect data object. The Encryption of bit stream is indicated in bit_31 of DWord_1. For such encrypted bit streams, this command has valid information in DWd_5, to load in the Initial counter value for AES-128 Counter mode decryption of the bit stream. DWd_1 Bit_30 is defined to indicate integrated encryption mode. For Bit streams which are integrated encrypted, this command has valid information in DWd_6, communicating the raw data length of bitstream and headers. Bit_29 is used to indicate the High Def. Video format. Before issuing an AVC_BSD_OBJECT command, all BSD states need to be valid. Therefore the commands used to set these states need to have been issued at some point prior to the issue of an AVC_BSD_OBJECT command.

VC1 deblock filter kernel cross the slice boundary if in the last MB row of a slice, so need to know the last MB row of a slice to disable the edge mask. There is why VC1 BSD H/W need to know the end of MB address for the current slice. As such no more phantom slice is needed for VC1, as long as the driver will program both start MB address in the current slice and the start MB address of the next slice. As a result, we can also support multiple picture state commands in between slices.

To handle encrypted bitstream decoding, the driver is required to align the 16byte chunk containing this first byte of actual bit-stream slice, to a naturally aligned 16byte boundary, i.e. that 16byte chunk should start on an aligned 16byte address boundary. The application should ensure that after the last valid byte of bitstream data and before encryption, enough padding is added up to the end of a 16byte chunk. In addition, for a protected context, all kernel source and destination data surface addresses must be in PCM space.

| Dword | Bits | Description |
|-------|-------|--|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | VC1 Command Opcode = VC1_BSD_OBJECT Pipeline[28:27] = BSD = 2h; Opcode[26:24] = VC1 = 5h; Sub Opcode[23:16] = 8h |
| | 15:0 | DWord Length (Excludes DWords 0,1) = 0007h |
| 1 | 31 | Bit stream is AES-128 Counter mode Encrypted '0' – Not encrypted '1' - Encrypted |
| | 28:22 | Reserved. MBZ |



| Dword | Bits | Description |
|-------|-------|---|
| | 21:0 | <p>Indirect Data's Actual Length of Valid Bitstream. This field provides the length in bytes of the indirect data. A value zero indicates that indirect data fetching is disabled – subsequently, the Indirect Data Start Address field is ignored.</p> <p>This field must have the same alignment as the Indirect Object Data Start Address.</p> <p>This field specifies the length in bytes of only the valid bit stream. This means that in the case of encrypted bit stream, the last byte of valid bit stream data, can line up on any byte aligned location within a 16byte encrypted chunk. This means that the app. will do the rounding or padding up to the end of the 16bytes chunk before encryption.</p> <p>22 bits support the maximum bitstream buffer size of (e.g. for 1080i) is $3072 \times 120 \times 68 / 8 = 3133440$ bytes, which meets the requirements of the VC1 Specification. It includes the byte that contains the First MB Bit Offset.</p> <p>.</p> <p>Format = U22 in bytes</p> |
| 2 | 31:29 | Reserved : MBZ |
| | 28:0 | <p>Indirect Data Start Address. This field specifies the Graphics Memory starting address of the data to be fetched into BSD Unit for processing. This pointer is relative to the MDE Indirect Object Base Address.</p> <p>For the VC1 for both encrypted and un-encrypted bitstream data, it is a byte aligned address. Note however for encrypted bit stream, the 16byte chunk containing this first byte of actual bit-stream slice, should be naturally aligned, i.e. that 16byte chunk should start on an aligned 16byte address boundary. In other words the first byte of that 16byte chunk should be aligned to a 16byte boundary. <i><Check with the VC1 design if they assume the starting byte of the actual bitstream is also 16byte aligned or it can be anywhere within the 16byte chunk></i></p> <p>.</p> <p>Range = [0 - 512MB]</p> <p>OPEN: Supporting concealment as 0 is not allowed here [Keep this open after we deal with error handling].</p> |
| 3 | 31:30 | Reserved : MBZ |
| | 31:24 | <p>SliceStartVertPos (Slice Vertical Position)</p> <p>This field specifies the position in y-direction of the first macroblock in the Slice in unit of macroblocks. For SecondField this value is reset to zero as oppoed to the VC1 spec Ref: 9.1.2 Slice Layer.</p> |
| | 23:3 | Reserved : MBZ |
| | 2:0 | <p>FirstMbBitOffset (First Macroblock Bit Offset)</p> <p>This field provides the bit offset of the first macroblock of the Slice in the first byte of the input compressed bitstream.</p> <p>Format: U3</p> |
| 4 | 31:30 | Reserved : MBZ |
| | 31:24 | <p>NextSliceStartVertPos (Next Slice Vertical Position)</p> <p>This field specifies the position in y-direction of the first macroblock in the next Slice in unit of macroblocks.</p> <p>This field is primarily used for error concealment. In the case that current slice is the last slice, this field should set to the height of picture (since y-direction is zero-based numbering)</p> |
| | 23:0 | Reserved : MBZ |



| Dword | Bits | Description |
|-------|-------|--|
| 5 | 31:0 | Initial counter value for AES-128 Counter mode decryption. 32-bit initial counter value (which does not need to be secure) for the AES counter mode decryption. This counter value should be applied in the AES Counter mode decryption of the first 16byte chunk of aligned encrypted slice data which contains the 1st starting byte of valid slice bitstream. If Bit stream is AES-128 Counter mode Encrypted is 0 , i.e. not encrypted bit stream, this field is ignored and should be MBZ. |
| 6 | 31:23 | Reserved. MBZ |
| | 22:0 | Indirect Data's Raw Data Length of Bitstream and Headers. This field provides the length in bytes of the indirect data. This DWord field is valid only for bitstream hw decryption. It includes the byte length from the beginning of the first byte of the bitstream till the last byte of the bitstream, including all the intermediate header bytes. In the case of encrypted bit stream, the last byte of valid bit stream data, can line up on any byte aligned location within a 16byte encrypted chunk. This means that the app. will do the rounding or padding up to the end of the 16bytes chunk before encryption. The length programmed has one extra bit than the actual byte length of the valid bit-stream data. Format = U23 in bytes |
| | 28:16 | Reserved : MBZ |
| 8 | 31:29 | Reserved : MBZ |

3.3 Context Switch

There is no pre-emption for the BCS pipeline; hence every command buffer is required to contain all the states setup (preamble). Specifically, CPU can not interrupt the BCS-BSD pipe, to stop the operation in the middle of decoding a bitstream data.

Switch of contexts can only be performed at picture boundary.

3.4 Error Handling and Error Concealment

Design objective is to keep the playback continue going forward in time even when errors are encountered (detected or not). Hence, the system will reconstruct and display the new picture as much as possible, without dropping the current picture or repeating the previous decoded picture. Another aspect of the design is to minimize the interaction/intervention of driver software when non-fatal error is encountered.

Error handling is performed only at the Slice boundary, and in between pictures.

To catch the case when the VC1_BSD has finished decoding the current picture but the last macroblock has not reached, the driver needs to specify the end MB Y address of the current picture. This is implemented by having the driver send out the NextSliceStartVert MB y-address as a marker for the end of current slice/picture.

Note that due to the nature of splitting the decoding pipeline into the front end BSD and the back end picture reconstruction, error concealment must be performed by the front end BSD. Otherwise, the commands generated for the back end picture reconstruction may be incomplete and therefore cause further problems (including hardware hang conditions).



3.4.1 Design Assumptions

Operating Assumptions:

- The BSD unit is either actively decoding a picture, or in between pictures when errors occur.
- At the beginning of decoding a new picture (the first slice of a picture), the current MB location registers are reset to 0 by the H/W automatically
- The end of a picture is tracked by comparing the current MB location (just being decoded) to the picture size parameters (in MB unit) in the BSD unit. The BSD unit will never attempt to decode beyond the picture size programmed, even when there are data remained at the BSD input.
- The end of a picture sequence is only known in the driver. BSD unit will continue operating until no more decoding request commands are injected or when there is an optional time-out by the driver.
- BSD unit has no notion of picture structure; it only deals with Slice as instructed and directed by the State Commands. It does not know it is decoding the last slice of a picture.
- There is no H/W Start_Code_Detector at the front end of the BSD unit, so the end of Slice signal can only be detected and generated internally.



3.4.2 Error Concealment and Recovery Strategies

Depending on the category and position of error being detected, a concealment strategy will be adopted.

- Picture Repeat
 - Handled by the display controller – when it does not receive a new picture to display at the appropriate time period
- H/W MB Fillings (as described below)
- S/W Re-sync at the next Slice or picture boundary
 - Dropping a Slice, or
 - The driver can ignore the rest of the bitstream until the start of the next NAL is detected. That is to start decoding again at slice boundary
 - Dropping a picture entirely.
 - The driver can decide to ignore all subsequent slices until a new IDR has arrived, and to start decoding again from there on.

3.4.3 Error Handling

3.4.3.1 S/W Driver

Provide high level error handling capabilities – interrupts and/or polling. Perform error handling and recovery in between Slices and Pictures boundaries.

It also checks for the bitstream syntax errors at and above Slice Header Layers, before proceeding with the Slice decoding.

Currently there is no communication of error handling at the application level.

There are also two error status bits generated in the H/W that the S/W can poll regularly:

- Force_In_Progress
 - A H/W Status Bit (read-only) to indicate filling missing MBs is in progress
- Error_In_Slice
 - A H/W Status Bit (read-only) to indicate an error has detected in the HW, including VLC Decoding Error, BSD Decoded Data Error and MPR Error



3.4.3.2 H/W MB Filling

This is the only error handling mechanism built-in the BSD unit. It will automatically perform error detection and error concealment in the Slice Data Layer at the Slice boundary. There is no direct S/W control.

When decoding a Slice (contains multiple MBs), internally the BSD unit will keep track of the number of MBs being decoded so far. This number also represents the position of the corresponding MB in a picture, assuming a raster-scan MB order (top to bottom, left to right). Hence, the Current_MB_Number register is maintained and updated in the H/W automatically.

For each Slice Data, there is also a Start_MB_Number presented at the corresponding Slice Header. This number specifies the picture location of the very first MB decoded from this Slice Data.

Same approach as AVC, except that VC1 also has the additional information for end of Slice detection. VC1 driver snoops the starting MB address of next slice and send this information as a state in the BSD media object command to process the current slice. Therefore, when processing the current slice, the VC1-BSD engine already knows the end MB address.

1. When MB address overflow - VC1 is not the same as AVC, as VC1 has the last MB address, so it will always stop right at the last MB of the current slice. The only thing needs to do is :
 - a. VC1 will flush out the extra data in the bitstream for the current slice by VIN
2. Premature ending of a slice
 - a. no rewind mechanism in VC1
 - b. compare the end of MB address to the start address of next slice
 - 1) always complete the current erroneous MB to the end even if there is error, and the error concealment will start in the next MB address, only if the error concealment is enabled; otherwise stop there, and sent out end of slice, and start decoding the next slice – so leaving a hole in the picture.
 1. need to double check if error concealment is disabled, will the hole causing problem in the VFE pipe ?? Yes it is a problem, must fill in.
 - a. There is no explicit internal looping through all MB in VFE pipe, everything is driven media object commands (either from the driver if it is in IT/MC mode, or from BSD in VLD mode). So, if some MB is missing, means only less media object commands being sent to the VFE.
 - b. Driver predefined the size of the memory record based on the number of MB in the current slice and set up the head and tail pointers; there is no way to modify the tail pointers if any MB is missing – hence the CS will hang or crash if there is missing MB record. Because CS will fetch an invalid command
 - i. In addition, for overlap smoothing in which there is neighbor dependency (no an issue in VC1 MC and ILDB), the setup of scoreboard depends on the total number of MB, it expects there are so many threads to be generated to cover the current slice. If a MB is missing, the scoreboard will wait forever and hang.



- c. Error concealment regardless of picture type - interlace, interlace frame or progressive frame
 - 2) Intra
 - 1. I-picture
 - 2. Only DC is predicted from neighbor, supply one DC coefficient per block and send out EOB. That is all AC coefficients are effectively set to 0.
 - a. The previous approach of forcing all coefficients including DC to 0 has removed
 - 3. DC coefficient is sent to IT for further processing
 - 3) Inter
 - 4. assume to be P-MB in both P and B-Slice
 - 5. 1MV Skip MB
 - 6. Default to frame MB and use frame reference motion for interlaced frame picture
 - 7. No IT, only MC (MV = 0,0, use the nearest reference frame)
- 3. There is no phantom slice being sent from driver at the end of each VC1 picture

Application will tell the driver the end of frame with the last slice. Driver will wait and collect all the slices for the whole frame before processing.

3.4.3.3 Error Statistics

TBD

3.4.3.4 WatchDog Timers

At the driver level an additional error detection mechanism is needed. In the situation where errors have stalled the BSD unit from progressing, there must be a way to tell that the H/W is hung. Watch-dog timers are preset with some known timing constraints that are related to the decoding, e.g. an expected decoding time of a slice or a picture. Once the threshold has reached, a stall condition is assumed, and action will take place to reset the BSD unit, and up to the driver to decide where to start decoding again or simply abort the decoding processing.

Bitplane is decoded in the HOST for now. If later if the performance is an issue, will design H/W to do it.

Software/Application prefers single call for a frame that has multiple slice, vs multiple iteration calls for each slice in a frame.



3.5 AVC Decoder Commands

3.5.1 MFD_AVC_BSD_OBJECT Command

The MFD_AVC_BSD_OBJECT command is the only primitive command for the AVC Decoding Pipeline. The same command is used for both CABAC and CAVLD modes. The Slice Data portion of the bitstream is loaded as indirect data object.

Before issuing a MFD_AVC_BSD_OBJECT command, all AVC states of the MFD Engine need to be valid. Therefore the commands used to set these states need to have been issued prior to the issue of a MFD_AVC_BSD_OBJECT command.

Context switch interrupt is not supported by this command.

| Dword | Bits | Description |
|-------|-------|--|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | Command Opcode = MFD_AVC_BSD_OBJECT Pipeline[28:27]=2h; Opcode[26:24] = AVC = 1h; SubOpA[23:21] = Dec = 1h; SubOpB[20:16] = 8h |
| | 15:12 | Reserved : MBZ |
| | 11:0 | DWord Length (Excludes DWords 0,1) = 0004h |
| 1 | 31:22 | Reserved. MBZ |
| | 21:0 | Indirect BSD Data Length. This field provides the length in bytes of the indirect data. A value zero indicates that indirect data fetching is disabled – subsequently, the Indirect Data Start Address field is ignored. This field must have the same alignment as the Indirect Object Data Start Address. It is the length in bytes of the bitstream data for the current slice. It includes the first byte of the first macroblock and the last byte of the last macroblock in the slice. Specifically, the zero-padding bytes (if present) and the next start-code are excluded. Hardware ignores the contents after the last non-zero byte. This field is sized to support AVC High Profile 4.1 Level bitstream. Interpreted from the AVC Spec, the maximum number of bits per macroblock for 4:2:0 is 3072 (bounded by IPCM (16x16+64x2)x8). So the maximum slice size (e.g. for 1080i) is 3072 x120 x 68 / 8 = 3133440 bytes, which requires 22 bits. It includes the byte that contains the First MB Bit Offset In implementing a phantom slice at the end of a picture for automatic error concealment, this field should set to 0. Format = U22 in bytes |
| 2 | 31:29 | Reserved : MBZ |



| Dword | Bits | Description |
|-------|--------------|--|
| | 28:0 | <p>Indirect BSD Data Start Address. This field specifies the Graphics Memory starting address of the data to be fetched into BSD Unit for processing. This pointer is relative to the MFD Indirect Object Base Address.</p> <p>Hardware ignores this field if indirect data is not present.</p> <p>It is a byte-aligned address for the AVC bitstream data in both CABAC/CAVLD Modes.</p> <p>In implementing a phantom slice at the end of a picture for automatic error concealment, this field should set to 0. Range = [0 - 512MB]</p> <p>Error Handling: For phantom slice or any slice with Indirect BSD Data Length equals 0, this field must program to 0 as well.</p> |
| 3-5 | 31:0 Each | <p>Inline Data</p> <p>All the required Slice Header parameters and error handling settings are captured as inline data of the AVC_BSD_OBJECT command. It has a fixed size of 4 DWs. Its definition is described in the next section.</p> |

3.5.1.1 Inline Data Description

The Inline Data includes all the required slice decoding states, extracted primarily from the Slice Header and its derivatives. It provides information for the following operations:

5. CABAC/CAVLD decoding
6. Internal error handling at the Slice boundary
 - b. Hardware Automatic Error Concealment
7. Motion vector prediction decoding (MPR)
8. BSD output compositing (feeding the subsequent IT/MC/ILDB operations)

The only hardware error detection and concealment mechanism is comparing the Slice_Start_MB_Num of decoding the new current slice with the Current_MB_Num resulted from decoding the previous slice. If they do not match, an error is assumed. There are two possible cases, either “greater than” or “less than”. If the Slice_Start_MB_Num is less than the Current_MB_Num, the Current_MB_Num is adjusted internally to be the same as the Slice_Start_MB_Num. If the Slice_Start_MB_Num is greater than, there is a gap between the Current_MB_Num and the Slice_Start_MB_Num that needs to be filled. The filling is started from [Current_MB_Num – Rewind_Num], and the method of filling can be either intra or inter as specified in the state descriptor.

These state/parameter values may subject to change on a Slice boundary, and must be provided in each AVC_BSD_OBJECT command. The values set for these variables are retained internally, until they are reset by hardware Asynchronous Reset or changed by the next AVC_BSD_OBJECT command.



| Dword | Bit | Description |
|-------|--|---|
| 3 | 31 | <p>Concealment Method</p> <p>This field specifies the method used for concealment when error is detected. If set, a copy from collocated macroblock location is performed from the concealment reference indicated by the ConCeal_Pic_Id field. If it is not set, a copy from the current picture is performed using Intra 16x16 Prediction method.</p> <p>0 – Intra 16x16 Prediction 1 – Inter P Copy</p> |
| | 30 | <p>Init Current_MB_Number</p> <p>When set, the current Slice_Start_MB_Num, Slice_MB_Start_Hor_Pos and Slice_MB_Start_Vert_Pos fields will be used to initialize the Current_MB_Number register. This effectively disables the concealment capability.</p> |
| | 29:28 | Reserved : MBZ |
| | 27:24 | <p>Rewind_Num</p> <p>This field provides the number of MBs or MBAFF pairs to rewind when performing concealment. This is ignored if Slice_Start_MB_Num is smaller than or equal to the Current_MB_Num.</p> <p>Error Concealment MB start position = Current_MB_NUMBER – MIN { Decoded_MB_NUMBER, (MBAFF ? 2 : 1)*Force_Skip_Rewind }</p> <p>Decoded_MB_NUMBER = Number of MBs decoded in the previous slice.</p> <p>No longer used</p> |
| | 23:22 | Reserved : MBZ |
| | 21:16 | <p>Conceal_Pic_Id (Concealment Picture ID)</p> <p>This field identifies the picture in the reference list to be used for concealment. This field is only valid if Concealment Method is Inter P Copy.</p> <p>Bit 21 = 0 – Frame Picture = 1 – Field Picture</p> <p>Bit 20:16 – Frame Store Index[4:0]</p> |
| | 15 | Reserved : MBZ |
| | 14 | <p>BSDPrematureComplete Error Handling</p> <p>1 – Set the interrupt to the driver (provide MMIO registers for MB address R/W) (???) 0 – Ignore the error and continue (masked the interrupt), assume the hardware automatically perform the error handling</p> <p>It occurs in situation where the Slice decode is completed but there are still data in the bitstream.</p> |
| | 13 | Reserved : MBZ |
| | 12 | <p>MPR Error (MV out of range) Handling– what to do when the specific error has occurred</p> <p>1 – Set the interrupt to the driver (provide MMIO registers for MB address R/W) (???) 0 – Ignore the error and continue (masked the interrupt), assume the hardware automatically perform the error handling</p> |
| 11 | Reserved : MBZ | |
| 10 | <p>Entropy Error Handling – what to do when the specific error has occurred</p> <p>1 – Set the interrupt to the driver (provide MMIO registers for MB address R/W) (???) 0 – Ignore the error and continue (masked the interrupt), assume the hardware automatically perform the error handling</p> | |



| Dword | Bit | Description |
|-------|-------|---|
| | 9 | Reserved : MBZ |
| | 8 | MB Header Error Handling – what to do when the specific error has occurred 1 – Set the interrupt to the driver (provide MMIO registers for MB address R/W) (???) 0 – Ignore the error and continue (masked the interrupt), assume the hardware automatically perform the error concealment. |
| | 7:0 | Reserved : MBZ |
| 4 | 31:16 | First_MB_Byte_Offset of Slice Data from Slice Header It gives the byte offset to locate the first MB data in the bitstream for a slice. It works with the Slice_Header_Skip_Mode to determine how to read and parse the data provided by the Indirect BSD Data Start Address. If Slice_Header_Skip_Mode is 0, First_MB_Byte_Offset must be programmed to 0. If Slice_Header_Skip_Mode is 1, First_MB_Byte_Offset provides a byte count offset, but it assumes all Emulation Prevention Bytes have already removed from the Slice Header Portion of the bitstream. For example, if there are 3 Emulation Prevention Byte in the Slice Header, the final offset to the Slice Data will be First_MB_Byte_Offset + 3. |
| | 15:8 | Reserved : MBZ (DW4 is defined differently for Encoder) |
| | 7 | Fix_Prev_Mb_Skipped Enables an alternative method for decoding mb_skipped, to cope with an encoder that codes a skipped MB as a direct MB with no coefficient. |
| | 6 | Slice_Header_Skip_Mode 0 – Driver passes to H/W only the Slice Data portion of the bitstream, with Emulation Prevention Byte straddling across Slice Header and Slice Data already removed (if any). In this mode, the First_MB_Byte_Offset must set to 0. 1 – Driver passes to H/W the entire Slice Header and Slice Data. Hence, the First_MB_Byte_Offset is used to locate the Slice Data from the beginning of the Slice Header position in the bitstream. Default is 0. |
| | 5:3 | Reserved : MBZ (DW4 is defined differently for Encoder) |
| | 2:0 | First_MB_Bit_Offset (First Macroblock Bit Offset) This field provides the bit offset of the first macroblock of the Slice in the first byte of the input compressed bitstream. Format: U3 |
| 5 | 31:0 | Reserved : MBZ |

Emulation Prevention Byte Removal is handled in 2 different ways that affects the definition of the Slice Data Buffer and H/W behavior. .

1. Application is required to remove any emulation prevention byte straddling across the Slice Header and Slice Data boundary in the bitstream
 - a. As such, application will pass to the driver the exact starting location of Slice Data (Byte Offset and Bit Offset) in the Slice buffer and guarantee there is no Emulation Prevention Byte at the beginning of Slice Data. So H/W does not need to do any before the Slice Data.



2. Application does not remove any emulation prevention byte straddling across the Slice Header and Slice Data boundary in the bitstream
 - a. As such application will pass to the driver the same set of information as above, but H/W is now required to scan the Emulation Prevention Byte from the beginning of Slice Header through to Slice Data, in order to locate the exact starting point of the Slice Data.
 - i. This is because the Slice Header Byte Offset does not include any emulation bytes found in the Slice Header.

3.6 AVC Encoder PAK Commands

Each PAK Commands is composed of a command op-code DW and one or more command data DWs (inline data). The size of each command is specified as part of the op-code DW. Most of the commands have fixed size, except some are allowed to be of variable length.

There is an inherent order of executing MFC PAK commands that driver must follow.



3.6.1 MFC_AVC_FQM_STATE Command

This is a frame-level state. Reciprocal Scaling Lists are always sent from the driver regardless whether they are specified by an application or the default/flat lists are being used. This is done to save the ROM (to store the default matrices) inside the PAK Subsystem. Hence, the driver is responsible for determining the final set of scaling lists to be used for encoding the current slice, based on the AVC Spec Table 7-2 (Fall-Back Rules A and B). For encoding, there is no need to send the `qm_list_flags[i]`, $i=0$ to 7 and `qm_present_flag` to the PAK, since Scaling Lists syntax elements are encoded above Slice Data Layer.

FQM Reciprocal Scaling Lists elements are 16-bit each, equal to $1/\text{ScaleValue}$. QM matrix elements are 8-bit each, equal to ScaleValue . However, in AVC spec., the Reciprocal Scaling Lists elements are not exactly equal to one-over of the corresponding Scaling Lists elements. The numbers are adjusted to simplify hardware implementation.

For all the description below, a scaling list set contains 6 4x4 scaling lists (or forward scaling lists) and 2 8x8 scaling lists (or forward scaling lists).

In MFX PAK mode, PAK needs both forward Q scaling lists and IQ scaling lists. The IQ scaling lists are sent as in MFD in raster scan order as shown in `MFX_AVC_QM_STATE`. But the Forward Q scaling lists are sent in transport form, i.e. column-wise raster order (column-by-column) to simplify the H/W.

Precisely, if the reciprocal forward scaling matrix is $F[4][4]$, then the 16 word of the matrix will be set as the following:

| Dword | Bit | Description |
|----------|----------|--|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | Command Opcode = MFC_AVC_FQM_STATE Pipeline[28:27]=2h; Opcode[26:24] = AVC = 1h; SubOpA[23:21] = Enc = 2h; SubOpB[20:16] = 2h |
| | 15:12 | Reserved : MBZ |
| | 11:0 | DWord Length (Excludes DWords 0,1) = 6Fh |
| 8 Dwords | 16 words | Luma4x4 Intra InvWeightScale (i=0) Unsigned integer value, ranging from 0 to 65535. Interpreted as a fixed-point fractional values, not necessary exactly equal to reciprocal of the corresponding WeightScale(i=0). An array of 4x4 scaling values (one-to-one correspondence with the coefficient position) for an Intra Luma block that is stored in raster order. It is set to the values derived from the syntax elements in the current active PPS and active SPS. Format for each element: U16, ranging from 0 to 65535. |
| 8 Dwords | 16 words | Cb4x4 Intra InvWeightScale (i=1) Format for each element: U16, ranging from 0 to 65535. |
| 8 Dwords | 16 words | Cr4x4 Intra InvWeightScale (i=2) Format for each element: U16, ranging from 0 to 65535. |



| Dword | Bit | Description |
|--------------|-------------|--|
| 8 Dwords | 16 words | Luma4x4 Inter InvWeightScale (i=3) Format for each element: U16, ranging from 0 to 65535. |
| 8 Dwords | 16 words | Cb4x4 Inter InvWeightScale (i=4) Format for each element: U16, ranging from 0 to 65535. |
| 8 Dwords | 16 words | Cr4x4 Inter InvWeightScale (i=5) Format for each element: U16, ranging from 0 to 65535. |
| 32 Dwords | 64 words | Luma8x8 Intra InvWeightScale (i=6) An array of 8x8 scaling values (one-to-one correspondence with the coefficient position) for an Intra Luma block that is stored in row major order (i.e. raster scan order). Format for each element: U16, ranging from 0 to 65535. |
| 32 Dwords | 64 words | Luma8x8 Inter InvWeightScale (i=7) An array of 8x8 scaling values (one-to-one correspondence with the coefficient position) for an Inter Luma block that is stored in row major order (i.e. raster scan order). Format for each element: U16, ranging from 0 to 65535. |

3.6.2 MFC_AVC_PAK_INSERT_OBJECT Command

The MFC_AVC_PAK_INSERT_OBJECT command is the first primitive command for the AVC Encoding Pipeline.

This command is issued

- to setup the control and parameters of inserting a chunk of compressed/encoded bits into the current bitstream output buffer starting at the specified bit location
- to perform the actual insertion by transferring the command inline data to the output buffer max. 32 bits at a time.

It is a variable length command as the data to be inserted are presented as inline data of this command. It is a multiple of 32-bit (1 DW), as the data bus to the bitstream buffer is 32-bit wide.

Multiple insertion commands can be issued back to back in a series. It is host software's responsibility to make sure their corresponding data will properly stitch together to form a valid H.264 bitstream. Internally, MFX hardware will keep track of the very last two bytes' (the very last byte can be a partial byte) values of the previous insertion. It is required that the next Insertion Object Command or the next PAK Object Command to perform the start code emulation sequence check and prevention 0x03 byte insertion with this end condition of the previous insertion.

Hardware will keep track of an output bitstream buffer current byte position and the associated next bit insertion position index.

Data to be inserted can be a valid H.264 NAL units or a partial NAL unit. Certain NAL unit has a minimum byte size requirement. As such the hardware will optionally (enabled by STATE Command) determine the number of CABAC_ZERO_WORD to be inserted to the end of the current NAL, based on the minimum byte size of a NAL and the actual bin count of the encoded Slice. Since prior to the CABAC_ZERO_WORD insertion, the RBSP or EBSP is already byte-aligned, so each CABAC_ZERO_WORD insertion is actually a 3-byte sequence 0x00 00 03.



The inline data may have already been processed for start code emulation byte insertion, except the possibility of the last 2 bytes plus the very last partial byte (if any). Hence, when hardware performing the concatenation of multiple consecutive insertion commands, or concatenation of an insertion command and a PAK object command, it must check and perform the necessary start code emulation byte insert at the junction.

The inline data is required to be byte aligned on the left (first transmitted bit order) and may or may not be byte aligned on the right (last transmitted bits). The command will specify the bit offset of the last valid DW

Each insertion state command defines a chunk of bits (compressed data) to be inserted at a specific location of the output compressed bitstream in the output buffer.

Depend on CABAC or CAVLC encoding mode (from Slice State), PAK Object Command is always ended in byte aligned output bitstream except for CABAC header insertion which is bit aligned. In the aligned cases, PAK will perform 0 filling in CAVLC mode, and 1 filling in CABAC mode.

Insertion data can include:

- any encoded syntax elements bit data before the encoded Slice Data (PAK Object Command) of the current Slice
 - SPS NAL
 - PPS NAL
 - SEI NAL
 - Other Non-Slice NAL
 - Leading_Zero_8_bits (as many bytes as there is)
 - Start Code Prefix
 - NAL Header Byte
 - Slice Header
- Any encoded syntax elements bit data after the encoded Slice Data (PAK Object Command) of the current Slice and prior to the next encoded Slice Data of the next Slice or prior to the end of the bistream, whichever comes first
 - Cabac_Zero_Word (if its number is determined by Host ???)
 - Trailing_Zero_8bits (as many bytes as there is)
 - Anything listed above before a Slice Data

Context switch interrupt is not supported by this command.

| Dword | Bit | Description |
|-------|-------|--|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | Command Opcode = MFC_AVC_PAK_INSERT_OBJECT Pipeline[28:27]=2h; Opcode[26:24] = AVC = 1h; SubOpA[23:21] = Enc = 2h; SubOpB[20:16] = 8h |
| | 15:12 | Reserved : MBZ |
| | 11:0 | Dword Length (Excludes Dwords 0,1) = Variable Length in DW |
| 1 | 31:18 | Reserved : MBZ |
| | 17:16 | DataByteOffset – SrcDataStartingByteOffset[1:0] Source Data Starting Byte Position within the very first inline DW. |
| | 15:14 | Reserved : MBZ |



| Dword | Bit | Description |
|-------|------|---|
| | 13:8 | <p>DataBitsInLastDW – SrCDataEndingBitInclusion[5:0] Source Data to be included in the very last inline DW. Follows the MSBit is the upper bit of each byte within the DW. The lower byte is actually processed first. For example, SrCDataEndingBitInclusion = 9, bit 7:0 and bit 15 are included as valid header data. Valid range = 1 to 32.</p> |
| | 7:4 | <p>SkipEmulByteCnt – Skip Emulation Byte Count Skip emulation check for number of starting bytes It can be programmed from 0 to 15 bytes. For example, to skip the start code that has already prefixed in the bitstream.</p> |
| | 3 | <p>EmulationFlag – EmulationByteBitsInsertEnable 1 : instruct the hardware to perform Start Code Prefix (0x 00 00 01/02/03/00) Search and Prevention Byte (0x 03) insertion on the insertion data of this command. It is required that hardware will handle a start code prefix crossing the boundary between 2 insertion commands, or an insertion command followed by a PAK Object command.</p> |
| | 2 | <p>LastHeaderFlag – LastSrcHeaderDataInsertCommandFlag To process a series of consecutive insertion commands, this flag (=1) indicates the current command is the last 'header' insertion in the series. In CABAC, hardware must perform the "1" insert for byte align for Slice Header before Slice Data comes in in the next PAK-OBJECT command. In CAVLC, hardware ignores this bit.</p> |
| | 1 | <p>EndOfSliceFlag – LastDstDataInsertCommandFlag No more insertion command and no more PAK-OBJECT command follows. Flush data out to memory</p> |
| | 0 | <p>BitstreamStartReset – ResetBitStreamStartingPos (OPEN: This bit is redundant, the control is already in the Slice State command) 1 : Reset the bitstream buffer insertion position to the bitstream buffer starting position. 0 : Insert the current command inline data starting at the current bitstream buffer insertion position</p> |
| 2 | 31:0 | <p>InsertDataPayload First DW of actual Data to be inserted to the output bitstream buffer.</p> |
| : | 31:0 | |
| N | 31:0 | Last DW of actual Data to be inserted to the output bitstream buffer. |



3.6.3 MFC_STITCH_OBJECT Command

The MFC_STITCH_OBJECT command is used when stitch-enabled is set to 1, while CodecSel and StandardSel are set to ENCODE and AVC, respectively.

This command is used, for example, to stitch multiple bitstreams to form a transport stream. The input bitstreams may be elementary streams stored in protected memory and the output transport stream may be AES encrypted.

It is a variable length command as the data to be inserted are presented as either inline data and/or indirect data of this command.

Multiple insertion commands can be issued back to back in a series. It is host software's responsibility to make sure their corresponding data will properly stitch together to form a valid output.

Hardware keeps track of an output bitstream buffer current byte position and the associated next bit insertion position index.

Context switch interrupt is not supported by this command. When encryption is on, the output is subject to the alignment (address and length) restriction. In order to support interrupt, it is software's responsibility to set up ARB_ON/OFF commands at the proper position to allow interrupt.

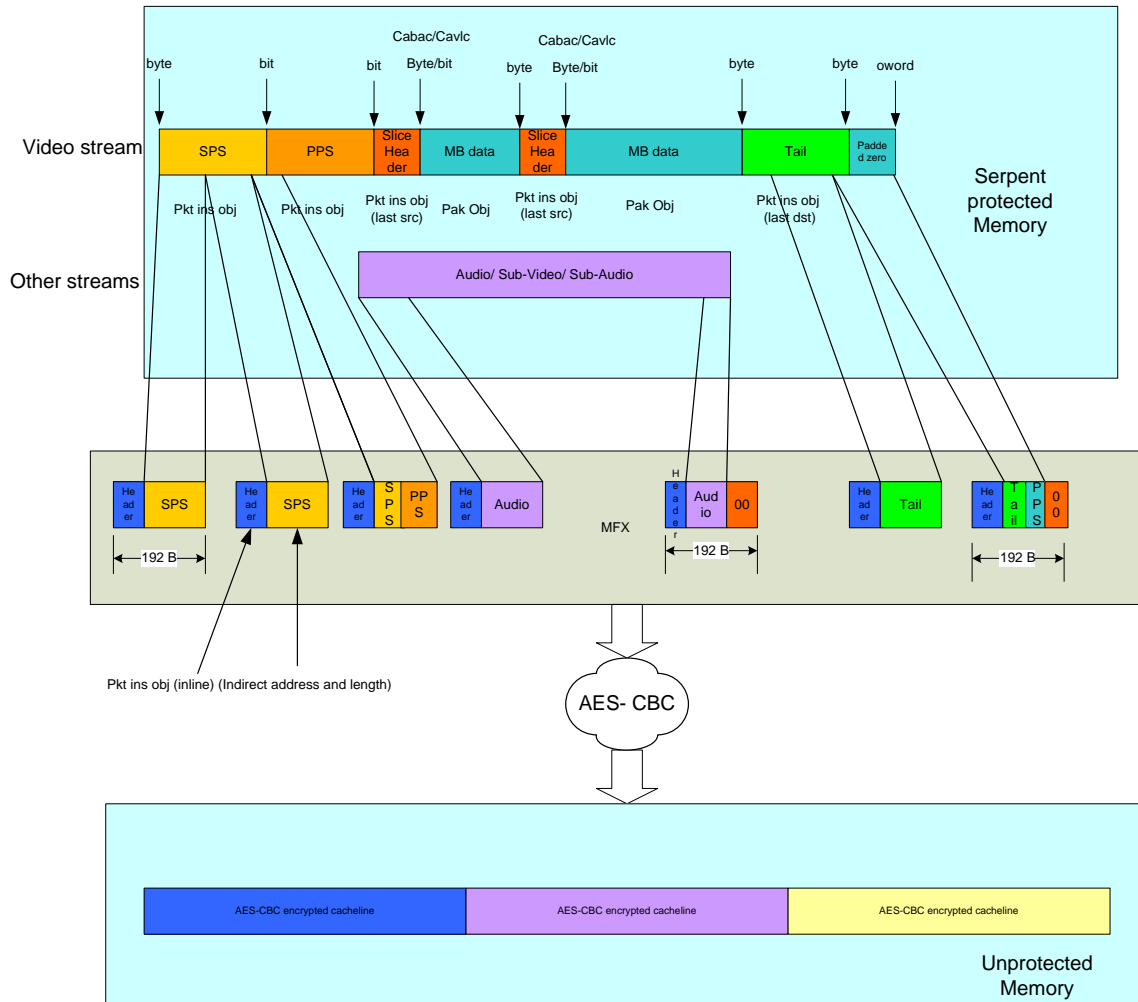
| Dword | Bit | Description |
|-------|---|---|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | Command Opcode = MFC_STITCH_OBJECT Pipeline[28:27]=2h; Opcode[26:24] = AVC = 1h; SubOpA[23:21] = Enc = 2h; SubOpB[20:16] = Ah |
| | 15:12 | Reserved : MBZ |
| | 11:0 | Dword Length (Excludes Dwords 0,1) = Variable Length in DW (>= 3) If it is 3, it indicates the absent of inline data. |
| 1 | 31:18 | Reserved : MBZ |
| | 17:16 | SrcDataStartingByteOffset[1:0] Source Data Starting Byte Position within the very first inline DW. |
| | 15:14 | Reserved : MBZ |
| | 13:8 | SrCDataEndingBitInclusion[5:0] Source Data to be included in the very last inline DW. Follows the MSBit is the upper bit of each byte within the DW. The lower byte is actually processed first. For example, SrCDataEndingBitInclusion =9, bit 7:0 and bit 15 are included as valid header data. Valid range = 1 to 32. |
| | 7:4 | EmulationSkipByteCount (DON'T CARE) Skip emulation check for number of starting bytes It can be programmed from 0 to 15 bytes. For example, to skip the start code that has already prefixed in the bitstream. |
| 3 | EmulationByteBitsInsertEnable (MUST BE ZERO) 1 : instruct the hardware to perform Start Code Prefix (0x 00 00 01/02/03/00) Search and Prevention Byte (0x 03) insertion on the insertion data of this command. It is required that hardware will handle a start code prefix crossing the boundary between 2 insertion commands, or an insertion command followed by a PAK Object command. | |



| Dword | Bit | Description |
|-------|-------|---|
| | 2 | LastSrcHeaderDataInsertCommandFlag To process a series of consecutive insertion commands, this flag (=1) indicates the current command is the last 'header' insertion in the series. In CABAC, hardware must perform the "1" insert for byte align for Slice Header before Slice Data comes in in the next PAK-OBJECT command. In CAVLC, hardware ignores this bit. |
| | 1 | LastDstDataInsertCommandFlag (THIS FIELD MUST BE THE SAME AS LastSrcHeaderDataInsertCommandFlag) No more insertion command and no more PAK-OBJECT command follows. Flush data out to memory |
| | 0 | ResetBitStreamStartingPos (Reserved: MBZ) 1 = Reset the bitstream buffer insertion position to the bitstream buffer starting position. 0 = Insert the current command inline data starting at the current bitstream buffer insertion position |
| 2 | 31:19 | Reserved : MBZ |
| | 18:0 | Indirect Data Length. This field provides the length in bytes of the indirect data. A value zero indicates that indirect data fetching is disabled – subsequently, the Indirect Data Start Address field is ignored. This field must have the same alignment as the Indirect Object Data Start Address. Format = U19 in bytes |
| 3 | 31:0 | Indirect Data Start Address. This field specifies the Graphics Memory starting address of the data to be loaded into the kernel for processing. This pointer is relative to the MFX Indirect Bitstream Object Base Address . Hardware ignores this field if indirect data is not present. Range = [0 – 4GB-1] |
| 4 | 31:0 | InsertDataPayLoad First DW of inline data to be inserted to the output bitstream buffer. |
| : | 31:0 | |
| N | 31:0 | Last DW of inline data to be inserted to the output bitstream buffer. |

3.6.3.1 Stitch Usage Example – AES Encrypted Transport Stream Output

The MFC_STITCH_OBJECT command may be used to stitch bitstreams (typically from several byte aligned input buffers) to form a transport stream with AES encryption. One usage example is depicted by the following picture. As the input buffers may be stored in the Serpent protected memory, host software (including graphics driver) doesn't have direct access to the data. Therefore, indirect data load is required for hardware to read the data and perform the transport packets generation. The input video elementary stream may be previously encoded (using MFC_AVC_PAK_OBJECT command) and assembled (using the MFC_AVC_PAK_INSERT_OBJECT command) through the MFX hardware. The transport packet header may contain small number of bytes. These data may be provided in the inline data of the MFC_STITCH_OBJECT commands.





3.6.4 MFC_AVC_PAK_OBJECT Command

The MFC_AVC_PAK_OBJECT command is the second primitive command for the AVC Encoding Pipeline. The same command is used for both CABAC and CAVLC modes. The MV Data portion of the bitstream is loaded as indirect data object.

Before issuing a MFC_AVC_PAK_OBJECT command, all AVC MFX states need to be valid. Therefore the commands used to set these states need to have been issued prior to the issue of this command.

MB record must be consecutive with no gaps, hence we do not need MB(x,y) in each MB command. Internal counter will keep track of the current MB address, starting from the Start_MB_In_Slice loaded at the beginning of each slice.

MFC_AVC_PAK_OBJECT command follows the MbType definition like MFD. Many fields in this command are identical to that in VME output. This is intended to reduce software converting overhead from VME to PAK.

Encoding statistical data such as the total size of the output bitstream are provided through MMIO registers. Software may access these registers through MI_STORE_REGISTER_MEM command.

| Dword | Bits | Description |
|-------|-------|---|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | Command Opcode = MFC_AVC_PAK_INSERT_OBJECT Pipeline[28:27]=2h; Opcode[26:24] = AVC = 1h; SubOpA[23:21] = Enc = 2h; SubOpB[20:16] = 9h |
| | 15:12 | Reserved : MBZ |
| | 11:0 | DWord Length (Excludes DWords 0,1) = 0009h |
| | 31:10 | Reserved. MBZ |
| 1 | 9:0 | Indirect PAK-MV Data Length This field provides the length in bytes of the indirect data, which contains all the MVs for the current MB (in any partitioning and subpartitioning form). A value zero indicates that indirect data fetching is disabled – subsequently, the Indirect PAK-MV Data Start Address field is ignored. This field must have the same alignment as the Indirect PAK-MV Data Start Address. This field must be DW aligned (since each MV is 4 bytes in size). Driver has to derived this field from MVsize *4 bytes per MV. Format = U10 in bytes |
| | 31:29 | Reserved : MBZ |
| 2 | 28:0 | Indirect PAK-MV Data Start Address Offset. This field specifies the memory starting address (offset) of the MV data to be fetched into PAK Subsystem for processing. This pointer is relative to the MFC Indirect PAK-MV Object Base Address . Hardware ignores this field if indirect data is not present, i.e. the Indirect PAK-MV Data Length is set to 0. It is a Dword aligned address in all AVC encoding configuration, since each MV is 4 bytes in size. Range = [0 - 512MB] |



| Dword | Bits | Description |
|-------|--------------|--|
| 3-10 | 31:0 Each | Inline Data All the required MB level controls and parameters for encoding are captured as inline data of the MFC_AVC_PAK_OBJECT command. It has a fixed size of 8 DWs. Its definition is described in the next section. |

3.6.4.1 PAK Object Inline Data Description

The Inline Data includes all the required MB encoding states, constitute part of the Slice Data syntax elements, MB Header syntax elements and their derivatives. It provides information for the following operations:

1. Forward and Inverse Transform
2. Forward and Inverse Quantization
3. Advanced Rate Control (QRC)
4. MB Parameter Construction (MPC)
5. CABAC/CAVLC encoding
6. Bit stream packing
7. Intra and inter-Prediction decoding loop
8. Internal error handling

These state/parameter values may subject to change on a per-MB basis, and must be provided in each MFC_AVC_PAK_OBJECT command. The values set for these variables are retained internally, until they are reset by hardware Asynchronous Reset or changed by the next MFC_AVC_PAK_OBJECT command.

The Deblocker Filter Control flags (FilterInternalEdgesFlag, FilterTopMbEdgeFlag and FilterLeftMbEdgesFlag) are generated by H/W, which are depending on MbaffFrameFlag, CurrMbAddr, PicWidthInMbs and disable_deblocking_filter_idc states.

Current MB [x,y] address is not sent, it is assumed that the H/W will keep track of the MB count and current MB position internally.



| DWord | Bit | Description |
|-------|---|---|
| 3 | 31:30 | Reserved: MBZ |
| | 23 | Reserved : MBZ (reserved for future use as ExternalMvBufFlag) |
| | 22:20 | <p>MvFormat (Motion Vector Size). This field specifies the size and format of the output motion vectors.</p> <p>This field is reserved (MBZ) when the IntraMbFlag = 1.</p> <p>The valid encodings are:</p> <p>000 = 0: No motion vector</p> <p>100 = 8MV: Four 8x8 motion vector pairs</p> <p>110 = 32MV: 16 4x4 motion vector pairs</p> <p>Others are reserved.</p> <p><i>(The following encodings are intended for future usages:</i></p> <p><i>001 = 1MV: one 16x16 motion vector</i></p> <p><i>010 = 2MV: One 16x16 motion vector pair</i></p> <p><i>011 = 4MV: Four 8x8 motion vectors</i></p> <p><i>101 = 16MV: 16 4x4 motion vectors</i></p> <p><i>111 = Packed, number of MVs is given by PackedMvNum.)</i></p> <p>Note:</p> |
| | 19 | <p>CbpDcY. This field specifies if the Luma DC sub-block is coded. Setting it to 0 will force PAK to zero out the Luma sub-block. Otherwise, whether the sub-block is coded will be determined by the quantization process.</p> <p>1 – the 4x4 DC-only Luma sub-block of the Intra16x16 coded MB is present; it is still possible that all DC coefficients are zero.</p> <p>0 – no 4x4 DC-only Luma sub-block is present; either not in Intra16x16 MB mode or all DC coefficients are zero.</p> |
| | 18 | <p>CbpDcU. This field specifies if the Chroma Cb DC sub-block is coded. Setting it to 0 will force PAK to zero out the Luma sub-block. Otherwise, whether the sub-block is coded will be determined by the quantization process.</p> <p>1 – the 2x2 DC-only Chroma Cb sub-block of all coded MB (any type) is present; it is still possible that all DC coefficients are zero.</p> <p>0 – no 2x2 DC-only Chroma Cb sub-block is present; all DC coefficients are zero.</p> |
| | 17 | <p>CbpDcV. This field specifies if the Chroma Cb DC sub-block is coded. Setting it to 0 will force PAK to zero out the Luma sub-block. Otherwise, whether the sub-block is coded will be determined by the quantization process.</p> <p>1 – the 2x2 DC-only Chroma Cr sub-block of all coded MB (any type) is present; it is still possible that all DC coefficients are zero.</p> <p>0 – no 2x2 DC-only Chroma Cr sub-block is present; all DC coefficients are zero.</p> |
| 16 | Reserved: MBZ (reserved for future use as ExternalResidBufFlag for turbo mode) | |

| DWord | Bit | Description |
|-------|------|--|
| | 15 | <p>Transform8x8Flag</p> <p>This field indicates that 8x8 transform is used for the macroblock.</p> <p>When it is set to 0, the current MB uses 4x4 transform. When it is set to 1, the current MB uses 8x8 transform. The transform_size_8x8_flag syntax element, if present in the output bitstream, is the same as this field. However, whether transform_size_8x8_flag is present or not in the output bitstream depends on several other conditions.</p> <p>This field is only allowed to be set to 1 for two conditions:</p> <ul style="list-style-type: none"> • It must be 1 if IntraMbFlag = INTRA and IntraMbMode = INTRA_8x8 • It may be 1 if IntraMbFlag = INTER and there is no sub partition size less than 8x8 <p>Otherwise, this field must be set to 0.</p> <p>0: 4x4 integer transform 1: 8x8 integer transform</p> |
| | 14 | <p>FieldMbFlag</p> <p>This field specifies the field polarity of the current macroblock, as the mb_field_decoding_flag syntax element in AVC spec.</p> <p>This field specifies whether current macroblock is coded as a field or frame macroblock in MBAFF mode. It is exactly the same as FIELD_PIC_FLAG syntax element in non-MBAFF mode.</p> <p>0 = Frame macroblock 1 = Field macroblock</p> |
| | 13 | <p>IntraMbFlag</p> <p>This field specifies whether the current macroblock is an Intra (I) macroblock. I_PCM is considered as Intra MB.</p> <p>For I-picture MB (IntraPicFlag =1), this field must be set to 1.</p> <p>This flag must be set in consistent with the interpretation of MbType (inter or intra modes).</p> <p>0: INTER (inter macroblock) 1: INTRA (intra macroblock)</p> |
| | 12:8 | <p>MbType5Bits</p> <p>This field is encoded to match with the best macroblock mode determined as described in the next section. It follows an unified encoding for inter and intra macroblocks according to AVC Spec.</p> |



| DWord | Bit | Description |
|-------|-------|--|
| | 7 | <p>FieldMbPolarityFlag</p> <p>This field indicates the field polarity of the current macroblock.</p> <p>Within an MbAff frame picture, this field may be different per macroblock and is set to 1 only for the second macroblock in a MbAff pair if FieldMbFlag is set. Otherwise, it is set to 0.</p> <p>Within a field picture, this field is set to 1 if the current picture is the bottom field picture. Otherwise, it is set to 0. It is a constant for the whole field picture.</p> <p>This field is reserved and MBZ for a progressive frame picture.</p> <p>0 = Current macroblock is a field macroblock from the top field</p> <p>1 = Current macroblock is a field macroblock from the bottom field</p> <p><i>Programming Note: Here bits [26:24] (MbAffFieldFlag and FiedlMbPolarityFlag) match with bits [10:8] of the Media Block Read message descriptor, simplifying the programming for message generation, as when MbAffFieldFlag is "1", kernels need to override the original "frame" surface state set for MBAFF frame picture.</i></p> |
| | 6 | Reserved: MBZ |
| | 5:4 | <p>IntraMbMode</p> <p>This field is provided to carry information partially overlapped with MbType.</p> <p>This field is only valid if IntraMbFlag = INTRA, otherwise, it is ignored by hardware..</p> |
| | 3 | Reserved: MBZ |
| | 2 | <p>SkipMbFlag</p> <p>By setting it to 1, this field forces an inter macroblock to be encoded as a skipped macroblock. It is equivalent to mb_skip_flag in AVS spec, indicating that a macroblock is inferred as a P_Skip (or B_Skip) in a P Slice (or B Slice). Hardware honors input MVs for motion prediction and forces CBP to zero.</p> <p>By setting it to 0, an inter macroblock will be coded as a normal inter macroblock. The macroblock may still be coded as a skipped macroblock, according to the macroblock type conversion rules described in the later sub sections.</p> <p>This field can only be set to 1 for certain values of MbType. See details later.</p> <p>This field is only valid for an inter macroblock. Hardware ignores this field for an intra macroblock.</p> <p>0 = not a skipped macroblock</p> <p>1 = is coded as a skipped macroblock</p> |
| | 1:0 | <p>InterMbMode</p> <p>This field is provided to carry redundant information as that encoded in MbType.</p> <p>This field is only valid if IntraMbFlag =0, otherwise, it is ignored by hardware.</p> |
| | 4 | 31:24 |
| | 23:16 | <p>MbYCnt (Vertical Origin). This field specifies the vertical origin of current macroblock in the destination picture in units of macroblocks.</p> <p>Format = U8 in unit of macroblock.</p> |
| | 15:8 | Reserved for future MbXCnt expansion. |



| DWord | Bit | Description | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|-------|-------|--|------|---|---|---|-------|----|----|----|------|---|---|---|-------|----|---|---|------|---|----|----|------|---|---|---|-------|----|----|----|------|---|---|---|------|---|------|---|------|---|------|---|
| | 31:16 | <p>[SNB only]: Cbp4x4Y[bit 15:0] (Coded Block Pattern Y) For 4x4 sub-block (when Transform8x8flag = 0 or in intra16x16) : 16-bit cbp, one bit for each 4x4 Luma sub-block (not including the DC 4x4 Luma block in intra16x16) in a MB. The 4x4 Luma sub-blocks are numbered as</p> <table border="0"> <tr> <td>blk0</td> <td>1</td> <td>4</td> <td>5</td> </tr> <tr> <td>bit15</td> <td>14</td> <td>11</td> <td>10</td> </tr> <tr> <td>blk2</td> <td>3</td> <td>6</td> <td>7</td> </tr> <tr> <td>bit13</td> <td>12</td> <td>9</td> <td>8</td> </tr> <tr> <td>blk8</td> <td>9</td> <td>12</td> <td>13</td> </tr> <tr> <td>bit7</td> <td>6</td> <td>3</td> <td>2</td> </tr> <tr> <td>blk10</td> <td>11</td> <td>14</td> <td>15</td> </tr> <tr> <td>bit5</td> <td>4</td> <td>1</td> <td>0</td> </tr> </table> <p>The cbpY bit assignment is cbpY bit [15 - X] for sub-block_num X.</p> <p>For 8x8 block (when Transform8x8flag = 1) Only the lower 4 bits [3:0] are valid; the remaining upper bits [15:4] are ignored. The 8x8 Luma blocks are numbered as</p> <table border="0"> <tr> <td>blk0</td> <td>1</td> <td>bit3</td> <td>2</td> </tr> <tr> <td>blk2</td> <td>3</td> <td>bit1</td> <td>0</td> </tr> </table> <p>The cbpY bit assignment is cbpY bit [3 - X] for block_num X. 0 in a bit - indicates the corresponding 8x8 block or 4x4 sub-block is not present (because all coefficient values are zero), or force to zero for PAK. 1 in a bit - indicates the corresponding 8x8 block or 4x4 sub-block is present (although it is still possible to have all its coefficients be zero - bad coding).</p> | blk0 | 1 | 4 | 5 | bit15 | 14 | 11 | 10 | blk2 | 3 | 6 | 7 | bit13 | 12 | 9 | 8 | blk8 | 9 | 12 | 13 | bit7 | 6 | 3 | 2 | blk10 | 11 | 14 | 15 | bit5 | 4 | 1 | 0 | blk0 | 1 | bit3 | 2 | blk2 | 3 | bit1 | 0 |
| blk0 | 1 | 4 | 5 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| bit15 | 14 | 11 | 10 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk2 | 3 | 6 | 7 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| bit13 | 12 | 9 | 8 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk8 | 9 | 12 | 13 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| bit7 | 6 | 3 | 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk10 | 11 | 14 | 15 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| bit5 | 4 | 1 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk0 | 1 | bit3 | 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk2 | 3 | bit1 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 15:8 | Reserved for future MbXCnt expansion. | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 15:8 | <p>[SNB only]: MbYCnt (Vertical Origin). This field specifies the vertical origin of current macroblock in the destination picture in units of macroblocks. Format = U8 in unit of macroblock.</p> | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 7:0 | <p>MbXCnt (Horizontal Origin). This field specifies the horizontal origin of current macroblock in the destination picture in units of macroblocks. Format = U8 in unit of macroblock.</p> | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 5 | 31:24 | <p>Reserved for future CbpAcUV expansion for 4.2.2. and 4.4.4 For 4.2.2, [23:16] for U(Cb), and [31:24] for C(Cr). For 4.4.4, the field [31:16] is interpreted as CbpAdj CbpAcV for 16 sub-blocks.</p> | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |



| DWord | Bit | Description | | | | | | | | |
|-------|-------|---|------|---|------|---|------|---|------|---|
| | 31:16 | <p>Cbp4x4V (Coded Block Pattern Cr)</p> <p>Only the lower 4 bits [3:0] are valid for 4:2:0. The 4x4 Cr sub-blocks are numbered as</p> <table style="margin-left: 40px;"> <tr> <td>blk0</td> <td>1</td> <td>bit3</td> <td>2</td> </tr> <tr> <td>blk2</td> <td>3</td> <td>bit1</td> <td>0</td> </tr> </table> <p>The cbpCr bit assignment is cbpCr bit [3 - X] for sub-block_num X.</p> <p>0 in a bit - indicates the corresponding 4x4 sub-block is not present (because all coefficient values are zero), or force to zero for PAK.</p> <p>1 in a bit - indicates the corresponding 4x4 sub-block is present (although it is still possible to have all its coefficients be zero - bad coding).</p> <p>For monochrome, this field is ignored. For SNB only 420 is supported.</p> <p>For 4.2.2, [23:16] for U(Cb), and [31:24] ignored.</p> <p>For 4.4.4, the definition is the same as for luma component: 1bit per 4x4 block..</p> | blk0 | 1 | bit3 | 2 | blk2 | 3 | bit1 | 0 |
| blk0 | 1 | bit3 | 2 | | | | | | | |
| blk2 | 3 | bit1 | 0 | | | | | | | |
| | 23:20 | <p>CbpAcV (Coded Block Pattern Cr)</p> <p>Only the lower 4 bits [3:0] are valid for 4:2:0. The 4x4 Cr sub-blocks are numbered as</p> <table style="margin-left: 40px;"> <tr> <td>blk0</td> <td>1</td> <td>bit3</td> <td>2</td> </tr> <tr> <td>blk2</td> <td>3</td> <td>bit1</td> <td>0</td> </tr> </table> <p>The cbpCr bit assignment is cbpCr bit [3 - X] for sub-block_num X.</p> <p>0 in a bit - indicates the corresponding 4x4 sub-block is not present (because all coefficient values are zero), or force to zero for PAK.</p> <p>1 in a bit - indicates the corresponding 4x4 sub-block is present (although it is still possible to have all its coefficients be zero - bad coding).</p> <p>For monochrome, this field is ignored. For SNB only 420 is supported.</p> | blk0 | 1 | bit3 | 2 | blk2 | 3 | bit1 | 0 |
| blk0 | 1 | bit3 | 2 | | | | | | | |
| blk2 | 3 | bit1 | 0 | | | | | | | |
| | 19:16 | <p>Cbp4x4U (Coded Block Pattern Cb)</p> <p>Only the lower 4 bits [3:0] are valid for 4:2:0. The 4x4 Cb sub-blocks are numbered as</p> <table style="margin-left: 40px;"> <tr> <td>blk0</td> <td>1</td> <td>bit3</td> <td>2</td> </tr> <tr> <td>blk2</td> <td>3</td> <td>bit1</td> <td>0</td> </tr> </table> <p>The cbpCb bit assignment is cbpCb bit [3 - X] for sub-block_num X.</p> <p>0 in a bit - indicates the corresponding 4x4 sub-block is not present (because all coefficient values are zero), or force to zero for PAK.</p> <p>1 in a bit - indicates the corresponding 4x4 sub-block is present (although it is still possible to have all its coefficients be zero - bad coding).</p> <p>For monochrome, this field is ignored. For SNB only 420 is supported.</p> | blk0 | 1 | bit3 | 2 | blk2 | 3 | bit1 | 0 |
| blk0 | 1 | bit3 | 2 | | | | | | | |
| blk2 | 3 | bit1 | 0 | | | | | | | |



| DWord | Bit | Description | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|-------|------|--|------|---|------|---|-------|----|------|----|------|---|---|---|-------|----|---|---|------|---|----|----|------|---|---|---|-------|----|----|----|------|---|---|---|------|---|------|---|------|---|------|---|
| | 15:0 | <p>Cbp4x4Y[bit 15:0] (Coded Block Pattern Y)</p> <p>For 4x4 sub-block (when Transform8x8flag = 0 or in intra16x16) :</p> <p>16-bit cbp, one bit for each 4x4 Luma sub-block (not including the DC block in intra16x16) in a MB. The 4x4 Luma sub-blocks are numbered as</p> <table border="0"> <tr> <td>blk0</td> <td>1</td> <td>4</td> <td>5</td> </tr> <tr> <td>bit15</td> <td>14</td> <td>11</td> <td>10</td> </tr> <tr> <td>blk2</td> <td>3</td> <td>6</td> <td>7</td> </tr> <tr> <td>bit13</td> <td>12</td> <td>9</td> <td>8</td> </tr> <tr> <td>blk8</td> <td>9</td> <td>12</td> <td>13</td> </tr> <tr> <td>bit7</td> <td>6</td> <td>3</td> <td>2</td> </tr> <tr> <td>blk10</td> <td>11</td> <td>14</td> <td>15</td> </tr> <tr> <td>bit5</td> <td>4</td> <td>1</td> <td>0</td> </tr> </table> <p>The cbpY bit assignment is cbpY bit [15 - X] for sub-block_num X.</p> <p>For 8x8 block (when Transform8x8flag = 1)</p> <p>Only the lower 4 bits [3:0] are valid; the remaining upper bits [15:4] are ignored. The 8x8 Luma blocks are numbered as</p> <table border="0"> <tr> <td>blk0</td> <td>1</td> <td>bit3</td> <td>2</td> </tr> <tr> <td>blk2</td> <td>3</td> <td>bit1</td> <td>0</td> </tr> </table> <p>The cbpY bit assignment is cbpY bit [3 - X] for block_num X.</p> <p>0 in a bit - indicates the corresponding 8x8 block or 4x4 sub-block is not present (because all coefficient values are zero), or force to zero for PAK.</p> <p>1 in a bit - indicates the corresponding 8x8 block or 4x4 sub-block is present (although it is still possible to have all its coefficients be zero - bad coding).</p> | blk0 | 1 | 4 | 5 | bit15 | 14 | 11 | 10 | blk2 | 3 | 6 | 7 | bit13 | 12 | 9 | 8 | blk8 | 9 | 12 | 13 | bit7 | 6 | 3 | 2 | blk10 | 11 | 14 | 15 | bit5 | 4 | 1 | 0 | blk0 | 1 | bit3 | 2 | blk2 | 3 | bit1 | 0 |
| blk0 | 1 | 4 | 5 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| bit15 | 14 | 11 | 10 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk2 | 3 | 6 | 7 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| bit13 | 12 | 9 | 8 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk8 | 9 | 12 | 13 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| bit7 | 6 | 3 | 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk10 | 11 | 14 | 15 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| bit5 | 4 | 1 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk0 | 1 | bit3 | 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk2 | 3 | bit1 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 15:0 | <p>[SNB only]:Cbp4x4U (Coded Block Pattern Cb)</p> <p>Only the lower 4 bits [3:0] are valid for 4:2:0. The 4x4 Cb sub-blocks are numbered as</p> <table border="0"> <tr> <td>blk0</td> <td>1</td> <td>bit3</td> <td>2</td> </tr> <tr> <td>blk2</td> <td>3</td> <td>bit1</td> <td>0</td> </tr> </table> <p>The cbpCb bit assignment is cbpCb bit [3 - X] for sub-block_num X.</p> <p>0 in a bit - indicates the corresponding 4x4 sub-block is not present (because all coefficient values are zero), or force to zero for PAK.</p> <p>1 in a bit - indicates the corresponding 4x4 sub-block is present (although it is still possible to have all its coefficients be zero - bad coding).</p> <p>For monochrome, this field is ignored. For SNB only 420 is supported.</p> <p>For 4.2.2, [7:0] for U(Cb), and [15:8] ignored.</p> <p>For 4.4.4, the definition is the same as for luma component: 1bit per 4x4 block.</p> | blk0 | 1 | bit3 | 2 | blk2 | 3 | bit1 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk0 | 1 | bit3 | 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk2 | 3 | bit1 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |



| DWord | Bit | Description |
|--------|--------------|--|
| 6 | 31:28 | <p>Skip8x8Pattern</p> <p>This field indicates whether each of the four 8x8 sub macroblocks is using the predicted MVs and will not be explicitly coded in the bitstream (the sub macroblock will be coded as direct mode). It contains four 1-bit subfields, corresponding to the 4 sub macroblocks in sequential order. The whole macroblock may be actually coded as B_Direct_16x16 or B_Skip, according to the macroblock type conversion rules described in a later sub section.</p> <p>This field is only valid for a B slice. It is ignored by hardware for a P slice. Hardware also ignores this field for an intra macroblock.</p> <p>0 in a bit – Corresponding MVs are sent in the bitstream 1 in a bit – Corresponding MVs are not sent in the bitstream</p> |
| | 27 | <p>EnableCoeffClamp</p> <p>1 = the magnitude of coefficients of the current MB will be clamped based on the clamping matrix after quantization 0 = no clamping</p> |
| | 26 | <p>LastMbFlag</p> <p>1 – the current MB is the last MB in the current Slice 0 – the current MB is not the last MB in the current SliceReserved MBZ.</p> |
| | 25 | <p>SkipMbConvDisable</p> <p>This is a per-MB level control to enable and disable skip conversion. This field is ORed with SkipConvDisable field. This field is only valid for a P or B slice. It must be zero for other slice types. Rules are provided in Section 3.6.4.1.6</p> <p>0 - Enable skip type conversion for the current macroblock 1 - Disable skip type conversion for the current macroblock</p> |
| | 24 | Reserved MBZ. |
| | 23:16 | <p>Reserved. Ignored by HW, this field will be re-derived internally. (was QpPrimeV. For 8-bit pixel data, QpCr is the same as QpPrimeCr, and it takes on a value in the range of 0 to 51, positive integer.)</p> |
| | 15:8 | <p>Reserved. Ignored by HW, this field will be re-derived internally. (Was QpPrimeU. For 8-bit pixel data, QpCb is the same as QpPrimeCb, and it takes on a value in the range of 0 to 51, positive integer.)</p> |
| | 7:0 | <p>QpPrimeY</p> <p>This is the per-MB QP value specified for the current MB. For 8-bit pixel data, QpY is the same as QpPrimeY, and it takes on a value in the range of 0 to 51, positive integer. Note: This value may differ from the actual codes, when HW QRC is on</p> |
| 7 to 9 | 31:0 Each | <p>For intra macroblocks, definition of these fields are specified in Error! Reference source not found. For inter macroblocks, definition of these fields are specified in Table 3-3</p> |
| 10 | 31:24 | <p>MaxSizeInWord</p> <p>PAK should not exceed this budget accumulatively, otherwise it will trickle the PANIC mode.</p> |
| | 23:16 | <p>TargetSizeInWord</p> <p>PAK should use this budget accumulatively to decide if it needs to limit the number of non-zero coefficients.</p> |



| DWord | Bit | Description |
|-------|------|----------------|
| | 15:0 | Reserved : MBZ |

The inline data content of Dwords 4 to 6 is defined either for intra prediction or for inter prediction, but not both.

Table 3-2. Inline data subfields for an Intra Macroblock

| Dword | Bit | Description |
|-------|-------|---|
| 7 | 31:16 | LumaIntraMode[1] Specifies the Luma Intra Prediction mode for four 4x4 sub-block of a MB, 4-bit each. See the bit assignment table later in this section. |
| | 15:0 | LumaIntraMode[0] Specifies the Luma Intra Prediction mode for four 4x4 sub-block, four 8x8 block or one intra16x16 of a MB. 4-bit per 4x4 sub-block (Transform8x8Flag=0, Mbtype=0 and intraMbFlag=1) or 8x8 block (Transform8x8Flag=1, Mbtype=0, MbFlag=1), since there are 9 intra modes. 4-bit for intra16x16 MB (Transform8x8Flag=0, Mbtype=1 to 24 and intraMbFlag=1), but only the LSBit[1:0] is valid, since there are only 4 intra modes. See the bit assignment table later in this section. |
| 8 | 31:16 | LumaIntraMode[3] Specifies the Luma Intra Prediction mode for four 4x4 sub-block of a MB, 4-bit each. See the bit assignment table later in this section. |
| | 15:0 | LumaIntraMode[2] Specifies the Luma Intra Prediction mode for four 4x4 sub-block of a MB, 4-bit each. See the bit assignment table later in this section. |
| 9 | 31:8 | Reserved : MBZ (Reserved for encoder turbo mode IntraResidueDataSize , when this is not 0, optional residue data are provided to the PAK; Reserved for decoder) |



| Dword | Bit | Description | | | | | | | | | | | | | | | | |
|-------|--|---|------|--------------------------------|---|---|---|---|---|--|---|---|---|---|---|---|-----|--|
| | 7:0 | <p>IntraStruct</p> <p>This field contains 6 bits for IntraPredAvailFlags[5:0] and 2 bits for ChromaIntraPredMode. The IntraPredAvailFlags[4:0] (the lower 5 bits) have already included the effect of the constrained_intra_pred_flag. See the diagram later for the definition of neighbor position around the current MB or MB pair (in MBAFF mode).</p> <p>1 – IntraPredAvailFlagY, indicates the values of samples of neighbor Y can be used in intra prediction for the current MB.</p> <p>0 – IntraPredAvailFlagY, indicates the values of samples of neighbor Y is not available for intra prediction of the current MB.</p> <p>IntraPredAvailFlag-A and -E can only be different from each other when constrained_intra_pred_flag is equal to 1 and mb_field_decoding_flag is equal to 1 and the value of the mb_field_decoding_flag for the macroblock pair to the left of the current macroblock is equal to 0 (which can only occur when MbaffFrameFlag is equal to 1).</p> <p>IntraPredAvailFlag-F is used only if</p> <ul style="list-style-type: none"> • it is in MBAFF mode, i.e. MbaffFrameFlag = 1, • the current macroblock is of frame type, i.e. MbFieldFag = 0, and • the current macroblock type is Intra8x8, i.e. IntraMbFlag = INTRA, IntraMbMode = INTRA_8x8, and Transform8x8Flag = 1. <p>In any other cases IntraPredAvailFlag-A shall be used instead.</p> <table border="1" data-bbox="550 982 1429 1371"> <thead> <tr> <th data-bbox="550 982 656 1045">Bits</th> <th data-bbox="656 982 1429 1045">IntraPredAvailFlags Definition</th> </tr> </thead> <tbody> <tr> <td data-bbox="550 1045 656 1087">7</td> <td data-bbox="656 1045 1429 1087">IntraPredAvailFlagF – F (Left 8th row (-1,7) neighbor)</td> </tr> <tr> <td data-bbox="550 1087 656 1129">6</td> <td data-bbox="656 1087 1429 1129">IntraPredAvailFlagA – A (Left neighbor top half)</td> </tr> <tr> <td data-bbox="550 1129 656 1171">5</td> <td data-bbox="656 1129 1429 1171">IntraPredAvailFlagE – E (Left neighbor bottom half)</td> </tr> <tr> <td data-bbox="550 1171 656 1213">4</td> <td data-bbox="656 1171 1429 1213">IntraPredAvailFlagB – B (Top neighbor)</td> </tr> <tr> <td data-bbox="550 1213 656 1255">3</td> <td data-bbox="656 1213 1429 1255">IntraPredAvailFlagC – C (Top right neighbor)</td> </tr> <tr> <td data-bbox="550 1255 656 1297">2</td> <td data-bbox="656 1255 1429 1297">IntraPredAvailFlagD – D (Top left corner neighbor)</td> </tr> <tr> <td data-bbox="550 1297 656 1371">1:0</td> <td data-bbox="656 1297 1429 1371">ChromaIntraPredMode – 2 bits to specify 1 of 4 chroma intra prediction modes, see the table in later section.</td> </tr> </tbody> </table> | Bits | IntraPredAvailFlags Definition | 7 | IntraPredAvailFlagF – F (Left 8 th row (-1,7) neighbor) | 6 | IntraPredAvailFlagA – A (Left neighbor top half) | 5 | IntraPredAvailFlagE – E (Left neighbor bottom half) | 4 | IntraPredAvailFlagB – B (Top neighbor) | 3 | IntraPredAvailFlagC – C (Top right neighbor) | 2 | IntraPredAvailFlagD – D (Top left corner neighbor) | 1:0 | ChromaIntraPredMode – 2 bits to specify 1 of 4 chroma intra prediction modes, see the table in later section. |
| Bits | IntraPredAvailFlags Definition | | | | | | | | | | | | | | | | | |
| 7 | IntraPredAvailFlagF – F (Left 8 th row (-1,7) neighbor) | | | | | | | | | | | | | | | | | |
| 6 | IntraPredAvailFlagA – A (Left neighbor top half) | | | | | | | | | | | | | | | | | |
| 5 | IntraPredAvailFlagE – E (Left neighbor bottom half) | | | | | | | | | | | | | | | | | |
| 4 | IntraPredAvailFlagB – B (Top neighbor) | | | | | | | | | | | | | | | | | |
| 3 | IntraPredAvailFlagC – C (Top right neighbor) | | | | | | | | | | | | | | | | | |
| 2 | IntraPredAvailFlagD – D (Top left corner neighbor) | | | | | | | | | | | | | | | | | |
| 1:0 | ChromaIntraPredMode – 2 bits to specify 1 of 4 chroma intra prediction modes, see the table in later section. | | | | | | | | | | | | | | | | | |



Table 3-3. Inline data subfields for an Inter Macroblock

| DWord | Bit | Description |
|-------|-------|--|
| 7 | 31:16 | Reserved : MBZ |
| | 15:8 | <p>SubMbPredMode (Sub-Macroblock Prediction Mode): If InterMbMode is INTER8x8, this field describes the prediction mode of the sub-partitions in the four 8x8 sub-macroblock. It contains four subfields each with 2-bits, corresponding to the four 8x8 sub-macroblocks in sequential order.</p> <p>This field is derived from sub_mb_type for a BP_8x8 macroblock.</p> <p>This field is derived from MbType for a non-BP_8x8 inter macroblock, and carries redundant information as MbType.</p> <p>If InterMbMode is INTER16x16, INTER16x8 or INTER8x16, this field carries the prediction modes of the sub macroblock (one 16x16, two 16x8 or two 8x16). The unused bits are set to zero.</p> <p>Bits [1:0]: SubMbPredMode[0] Bits [3:2]: SubMbPredMode[1] Bits [5:4]: SubMbPredMode[2] Bits [7:6]: SubMbPredMode[3]</p> |
| 8 | 31:24 | <p>RefPicSelect[0][3] up to 4 reference pictures per L0 direction, one per MB partition, if exists. See details in later section. This field specifies the reference index into the Reference Picture List0 Table.</p> |
| | 23:16 | <p>RefPicSelect[0][2] Support up to 4 reference pictures per L0 direction, one per MB partition, if exists. See details in later section. This field specifies the reference index into the Reference Picture List0 Table.</p> |
| | 15:8 | <p>RefPicSelect[0][1] Support up to 4 reference pictures per L0 direction, one per MB partition, if exists. See details in later section. This field specifies the reference index into the Reference Picture List0 Table.</p> |
| | 7:0 | <p>RefPicSelect[0][0] Support up to 4 reference pictures per L0 direction, one per MB partition, if exists. See details in later section. This field specifies the reference index into the Reference Picture List0 Table.</p> |
| 9 | 31:24 | <p>RefPicSelect[1][3] Support up to 4 reference pictures per L1 direction, one per MB partition, if exists. See details in later section. This field specifies the reference index into the Reference Picture List1 Table.</p> |
| | 23:16 | <p>RefPicSelect[1][2] Support up to 4 reference pictures per L1 direction, one per MB partition, if exists. See details in later section. This field specifies the reference index into the Reference Picture List1 Table.</p> |
| | 15:8 | <p>RefPicSelect[1][1] Support up to 4 reference pictures per L1 direction, one per MB partition, if exists. See details in later section. This field specifies the reference index into the Reference Picture List1 Table.</p> |



| DWord | Bit | Description |
|-------|-----|---|
| | 7:0 | RefPicSelect[1][0] Support up to 4 reference pictures per L1 direction, one per MB partition, if exists. See details in later section. This field specifies the reference index into the Reference Picture List1 Table. |

3.6.4.1.1 Luma Intra Prediction Modes

Luma Intra Prediction Modes (LumaIntraPredModes) are defined in Table 3-4. It is further categorized as Intra16x16PredMode (Table 3-5), Intra8x8PredMode (



Table 3-6) and Intra4x4PredMode (Table 3-7**Error! Reference source not found.**), operating on 16x16, 8x8 and 4x4 block sizes, respectively. Figure 3-1. Intra_4x4 prediction mode directions. Figure 3-1 illustrates the intra prediction directions geometrically for the Intra4x4 prediction. When a macroblock is subdivided, the intra prediction is performed for the subdivision in a predetermined order. For example, Figure 3-2. Numbers of Block4x4 in a 16x16 region shows the block order for Intra4x4 prediction. And Figure 3-3. Numbers of Block4x4 in an 8x8 region or numbers of Block8x8 in a 16x16 region shows the block order of Block8x8 in a 16x16 region or Block4x4 in an 8x8 region.

Table 3-4. Definition of LumaIntraPredModes

| LumaIntraPredModes [index] | | Intra16x16PredMode | Intra8x8PredMode | Intra4x4PredMode |
|-------------------------------|-------|---|------------------------------------|------------------------------------|
| Index | Bit | MbType = [1...24] Transform8x8Flag = 0 | MbType = 0 Transform8x8Flag = 1 | MbType = 0 Transform8x8Flag = 0 |
| 0 | 15:12 | MBZ | Block8x8 3 | Block4x4 3 (0_0) |
| | 11:8 | MBZ | Block8x8 2 | Block4x4 2 (0_1) |
| | 7:4 | MBZ | Block8x8 1 | Block4x4 1 (0_2) |
| | 3:0 | Block16x16 | Block8x8 0 | Block4x4 0 (0_3) |
| 1 | 15:12 | MBZ | MBZ | Block4x4 7 (1_0) |
| | 11:8 | MBZ | MBZ | Block4x4 6 (1_1) |
| | 7:4 | MBZ | MBZ | Block4x4 5 (1_2) |
| | 3:0 | MBZ | MBZ | Block4x4 4 (1_3) |
| 2 | 15:12 | MBZ | MBZ | Block4x4 11 (2_0) |
| | 11:8 | MBZ | MBZ | Block4x4 10 (2_1) |
| | 7:4 | MBZ | MBZ | Block4x4 9 (2_2) |
| | 3:0 | MBZ | MBZ | Block4x4 8 (2_3) |
| 3 | 15:12 | MBZ | MBZ | Block4x4 15 (3_0) |
| | 11:8 | MBZ | MBZ | Block4x4 14 (3_1) |
| | 7:4 | MBZ | MBZ | Block4x4 13 (3_2) |
| | 3:0 | MBZ | MBZ | Block4x4 12 (3_3) |



Table 3-5. Definition of Intra16x16PredMode

| Intra16x16PredMode | Description |
|---------------------------|-------------------------------|
| 0 | Intra_16x16_Vertical |
| 1 | Intra_16x16_Horizontal |
| 2 | Intra_16x16_DC |
| 3 | Intra_16x16_Plane |
| 4 – 15 | Reserved |



Table 3-6. Definition of Intra8x8PredMode

| Intra8x8PredMode | Description |
|-------------------------|-------------------------------|
| 0 | Intra_8x8_Vertical |
| 1 | Intra_8x8_Horizontal |
| 2 | Intra_8x8_DC |
| 3 | Intra_8x8_Diagonal_Down_Left |
| 4 | Intra_8x8_Diagonal_Down_Right |
| 5 | Intra_8x8_Vertical_Right |
| 6 | Intra_8x8_Horizontal_Down |
| 7 | Intra_8x8_Vertical_Left |
| 8 | Intra_8x8_Horizontal_Up |
| 9 – 15 | Reserved |

Table 3-7. Definition of Intra4x4PredMode

| Intra4x4PredMode | Description |
|-------------------------|-------------------------------|
| 0 | Intra_4x4_Vertical |
| 1 | Intra_4x4_Horizontal |
| 2 | Intra_4x4_DC |
| 3 | Intra_4x4_Diagonal_Down_Left |
| 4 | Intra_4x4_Diagonal_Down_Right |
| 5 | Intra_4x4_Vertical_Right |
| 6 | Intra_4x4_Horizontal_Down |
| 7 | Intra_4x4_Vertical_Left |
| 8 | Intra_4x4_Horizontal_Up |
| 9 – 15 | Reserved |



Figure 3-1. Intra_4x4 prediction mode directions




Figure 3-2. Numbers of Block4x4 in a 16x16 region

| | | | |
|----|----|----|----|
| 0 | 1 | 4 | 5 |
| 2 | 3 | 6 | 7 |
| 8 | 9 | 12 | 13 |
| 10 | 11 | 14 | 15 |



Figure 3-3. Numbers of Block4x4 in an 8x8 region or numbers of Block8x8 in a 16x16 region

| | |
|---|---|
| 0 | 1 |
| 2 | 3 |

Table 3-8. Definition of Chroma Intra Prediction Mode

| ChromaIntraPredMode (intra_chroma_pred_mode) | Name of intra_chroma_pred_mode |
|---|---|
| 0 | Intra_Chroma_DC (prediction mode) |
| 1 | Intra_Chroma_Horizontal (prediction mode) |
| 2 | Intra_Chroma_Vertical (prediction mode) |
| 3 | Intra_Chroma_Plane (prediction mode) |

3.6.4.1.2 Reference Indices defined for each MB partition type and Bit Assignment

| MB partitioning | frame/field MB/Picture | | | | |
|--------------------|------------------------|------|------|------|-----------|
| | 16x16 | 16x8 | 8x16 | 8x8 | |
| RefIdxL0/1[0] | blk0 | blk0 | blk0 | blk0 | Bit 7:0 |
| RefIdxL0/1[1] | x | blk1 | blk1 | blk1 | Bit 15:8 |
| RefIdxL0/1[2] | x | x | x | blk2 | Bit 23:16 |
| RefIdxL0/1[3] | x | x | x | blk3 | Bit 31:24 |



3.6.4.1.3 MB Neighbor Availability in Intra-Prediction Modes (IntraPredAvailFlags)

Current MB is labelled as X. For non-MBAFF mode, 4 neighbors, A, B, C, D, are depicted in the following picture and are defined as the following.

- MB D: top left neighbor of current MB X
- MB C: top right neighbor of current MB X
- MB B: top neighbor of current MB X
- MB A: left neighbor of the current MB X

| | | |
|----------------------------|-----------------------|-----------------------------|
| mbAddrD D (top-left) | mbAddrB B (top) | mbAddrC C (top-right) |
| mbAddrA A (left) | X CurrMbAddrX | N/A |
| N/A | N/A | N/A |

For MBAFF mode, the current MB is labelled as X0 or X1, 4 neighbor pairs, A0/A1, B0/B1, C0/C1, D0/D1, are depicted in the following picture and are defined as the following.

- MB D0: first MB of top left neighbor MB pair of current MB pair X0/X1
- MB D1: second MB of top left neighbor MB pair of current MB pair X0/X1
- MB C0: first MB of top right neighbor MB pair of current MB pair X0/X1
- MB C1: second MB of top right neighbor MB pair of current MB pair X0/X1
- MB B0: first MB of top neighbor MB pair of current MB pair X0/X1
- MB B1: second MB of top neighbor MB pair of current MB pair X0/X1
- MB A0: first MB of left neighbor MB pair of the current MB pair X0/X1
- MB A1: second MB of left neighbor MB pair of the current MB pair X0/X1

| | | |
|-----------------|-----------------|-----------------|
| mbAddrD D0 | mbAddrB B0 | mbAddrC C0 |
| mbAddrD+1 D1 | mbAddrB+1 B1 | mbAddrC+1 C1 |



| | | |
|-----------------|-------------------------|-----|
| mbAddrA A0 | CurrMbAddrX X0 or | N/A |
| mbAddrA+1 A1 | CurrMbAddrX X1 | N/A |

For a given macroblock X (or X0/X1), the 6 neighbor availability signals, namely, A, B, C, D, E, F, are defined as the following.

- IntraPredAvailFlagF – F: (Single neighbor pixel at the left 8th row (-1,7))
- IntraPredAvailFlagA – A (Left neighbor top half pixel group)
- IntraPredAvailFlagE – E (Left neighbor bottom half pixel group)
- IntraPredAvailFlagB – B (Top neighbor pixel group)
- IntraPredAvailFlagC – C (Top right neighbor pixel group)
- IntraPredAvailFlagD – D (Top left corner neighbor pixel)

The following table depicts the generation of IntraPredAvailFlags[5:0] signals in a condensed form. It should note that for most cases only one input neighbor signal is assigned for each condition. The exception is in the four places for deriving left neighbor A and E where the neighbor is only available if left neighbors (A0 and A1) are both available (A0&A1). Also note that F takes output value very similar to that for A except the two “AND” conditions, where F is assigned to A1 instead of (A0&A1).

Table 3-9. Definition of intra-prediction neighbor availability calculation in MBAFF mode

| Output → | | D | | B | | C | | A | | E | | F | |
|----------------------------|---------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|---------------------------------|----------------|---------------------------------|----------------|----------------|
| Current X \ Neighbor Y | | Y-Frame | Y-Field | Y-Frame | Y-Field | Y-Frame | Y-Field | Y-Frame | Y-Field | Y-Frame | Y-Field | Y-Frame | Y-Field |
| X ₀ (Top) | X-Frame | D ₁ | D ₁ | B ₁ | B ₁ | C ₁ | C ₁ | A ₀ | A ₀ & A ₁ | A ₀ | A ₀ & A ₁ | A ₀ | A ₁ |
| | X-Field | D ₁ | D ₀ | B ₁ | B ₀ | C ₁ | C ₀ | A ₀ | A ₀ | A ₁ | A ₀ | A ₀ | A ₀ |
| X ₁ (Bottom) | X-Frame | A ₀ | A ₁ | X ₀ | N/A | 0 | 0 | A ₁ | A ₀ & A ₁ | A ₁ | A ₀ & A ₁ | A ₁ | A ₁ |
| | X-Field | D ₁ | D ₁ | B ₁ | B ₁ | C ₁ | C ₁ | A ₀ | A ₁ | A ₁ | A ₁ | A ₀ | A ₁ |

In Table 3-9, *X-Frame* or *X-Field* indicates the frame/field mode of the current MB; and *Y-Frame* or *Y-Field* indicates the corresponding neighbor MB for the given neighbor location, being upper left (D) or left (A) for example. Therefore, “Y-” takes the selected neighbor MB name as in the output cell in the same column. For example, for output D, if X1 is a frame MB, Y = A, if X1 is a field MB, Y = D.



For non-MBAFF mode, as A0=A1, B0=B1, C0=C1 and D0=D1, the neighbor assignment is degenerated into the following simple table. Here, E is assigned to the same as A and F is forced to 0.

Table 3-10. Definition of intra-prediction neighbor availability calculation in non-MBAFF mode

| Output → | D | B | C | A | E | F |
|----------|----|----|----|----|----|---|
| X | D0 | B0 | C0 | A0 | A0 | 0 |

To further explain the neighbor assignment rules in Table 3-9, the following table provides description for each condition. Please note that this table is **informative** as it provides redundant information as in Table 3-9.

Table 3-11. Detailed explanation of intra-prediction neighbor availability calculation in MBAFF mode

| Current MB | Current MB Field | Neighbor MB Field | Neighbor MB Select (Y=?) | Neighbor Avail Result (OUTPUT) | Description |
|----------------|------------------|-------------------|--------------------------|--------------------------------|---|
| D | | | | | |
| X0 (Top) | X-Frame | Y-Frame | D | D1 | Top Frame MB uses [-1,-1] = D_31, thus D1 only, regardless D frame or field pair |
| | X-Frame | Y-Field | D | D1 | |
| | X-Field | Y-Frame | D | D1 | Top Field MB uses [-1,-2] = D_30, thus if D is frame pair, takes D1 (D1_14 pixel), and if D is field pair, takes D0 (D0_15 pixel) |
| | X-Field | Y-Field | D | D0 | |
| X1 (Bottom) | X-Frame | Y-Frame | A | A0 | Bottom Frame MB uses [-1,15] = A_15, thus A0 (A0_15 pixel) if A is a frame pair, or A1 (A1_7 pixel), if A is a field pair |
| | X-Frame | Y-Field | A | A1 | |
| | X-Field | Y-Frame | D | D1 | Bottom Field MB uses [-1,-1] = D_31, thus D1 only, regardless D frame or field pair |
| | X-Field | Y-Field | D | D1 | |
| B | | | | | |
| X0 (Top) | X-Frame | Y-Frame | B | B1 | Top Frame MB uses [0...15,-1] = B_31, thus B1 only, regardless B frame or field pair |
| | X-Frame | Y-Field | B | B1 | |
| | X-Field | Y-Frame | B | B1 | Top Field MB uses [0...15,-2] = B_30, thus if B is frame pair, takes B1 (B1_14 row), and if B is field pair, takes B0 (B0_15 row) |
| | X-Field | Y-Field | B | B0 | |
| X1 (Bottom) | X-Frame | Y-Frame | X | X0 | Bottom Frame MB uses [0...15,15], thus X0 (X0_15 row) |
| | X-Frame | Y-Field | X | n/a | <i>Note: X0 and X1 must have the same field type, this row is n/a.</i> |
| | X-Field | Y-Frame | B | B1 | Bottom Field MB uses [0...15,-1] = B_31, thus B1 only, regardless B frame or field pair |
| | X-Field | Y-Field | B | B1 | |
| C | | | | | |



| Current MB | Current MB Field | Neighbor MB Field | Neighbor MB Select (Y=?) | Neighbor Avail Result (OUTPUT) | Description |
|----------------|------------------|-------------------|--------------------------|--------------------------------|---|
| X0 (Top) | X-Frame | Y-Frame | C | C1 | Top Frame MB uses [16...23,-1] = C_31, thus C1 only, regardless C frame or field pair |
| | X-Frame | Y-Field | C | C1 | |
| | X-Field | Y-Frame | C | C1 | Top Field MB uses [16...23,-2] = C_30, thus if C is frame pair, takes C1 (C1_14 row), and if C is field pair, takes C0 (C0_15 row) |
| | X-Field | Y-Field | C | C0 | |
| X1 (Bottom) | X-Frame | Y-Frame | n/a | 0 | Bottom Frame MB doesn't have left-top neighbor by definition, thus forced to 0 |
| | X-Frame | Y-Field | n/a | 0 | |
| | X-Field | Y-Frame | C | C1 | Bottom Field MB uses [16...23,-1] = C_31, thus C1 only, regardless C frame or field pair |
| | X-Field | Y-Field | C | C1 | |
| | | | | A | |
| X0 (Top) | X-Frame | Y-Frame | A | A0 | First Half of Top Frame MB uses [-1,0...7], thus A0 if A is a frame pair; but is only avail if both A0 and A1 are avail if A is a field pair due to the mix |
| | X-Frame | Y-Field | A | A0&A1 | |
| | X-Field | Y-Frame | A | A0 | First Half of Top Field MB uses [-1,0..2..4..14], thus take A0 (if A is frame pair, takes A0 even lines, and if A is field pair, takes A0 first half) |
| | X-Field | Y-Field | A | A0 | |
| X1 (Bottom) | X-Frame | Y-Frame | A | A1 | First Half of Bottom Frame MB uses [-1,16...23], thus A1 if A is a frame pair; but is only avail if both A0 and A1 are avail if A is a field pair due to the mix |
| | X-Frame | Y-Field | A | A0&A1 | |
| | X-Field | Y-Frame | A | A0 | First Half of Bottom Field MB uses [-1,1..3..15], thus take A0 (if A is frame pair, takes A0 odd lines, and if A is field pair, takes A1 first half) |
| | X-Field | Y-Field | A | A1 | |
| | | | | E | |
| X0 (Top) | X-Frame | Y-Frame | A | A0 | Second Half of Top Frame MB uses [-1,8...15], thus A0 if A is a frame pair; but is only avail if both A0 and A1 are avail if A is a field pair due to the mix |
| | X-Frame | Y-Field | A | A0&A1 | |
| | X-Field | Y-Frame | A | A1 | Second Half of Top Field MB uses [-1,16..18..30], thus take A1 (if A is frame pair, takes A1 even lines, and if A is field pair, takes A0 second half) |
| | X-Field | Y-Field | A | A0 | |
| X1 (Bottom) | X-Frame | Y-Frame | A | A1 | Second Half of Bottom Frame MB uses [-1,24...31], thus A1 if A is a frame pair; but is only avail if both A0 and A1 are avail if A is a field pair due to the mix |
| | X-Frame | Y-Field | A | A0&A1 | |
| | X-Field | Y-Frame | A | A1 | Second Half of Bottom Field MB uses [-1,17..19..31], thus takes A1 (if A is frame pair, takes A1 odd lines, and if A is field pair, takes A1 second half) |
| | X-Field | Y-Field | A | A1 | |
| | | | | F | |
| X0 (Top) | X-Frame | Y-Frame | A | A0 | Top Frame MB uses [-1,7] = A_7 (odd location), thus A0 if A is frame pair and A1 if field pair |
| | X-Frame | Y-Field | A | A1 | |



| Current MB | Current MB Field | Neighbor MB Field | Neighbor MB Select (Y=?) | Neighbor Avail Result (OUTPUT) | Description |
|----------------|------------------|-------------------|--------------------------|--------------------------------|--|
| | X-Field | Y-Frame | A | A0 | Top Field MB uses [-1,14] = A_14 (even location), thus A0 regardless A frame or field pair |
| | X-Field | Y-Field | A | A0 | |
| X1 (Bottom) | X-Frame | Y-Frame | A | A1 | Bottom Frame MB uses [-1,23] = A_23 (odd location), thus A1 regardless A frame or field pair |
| | X-Frame | Y-Field | A | A1 | |
| | X-Field | Y-Frame | A | A0 | Bottom Field MB uses [-1,15] = A_15 (odd location), thus A0 if A is frame pair and A1 if A is field pair |
| | X-Field | Y-Field | A | A1 | |

3.6.4.1.4 Macroblock Type for Intra Cases

MbType follows two different tables according to whether the macroblock is an inter or intra macroblock according to IntraMbFlag.

For an intra macroblock, MbType, as defined in Table 3-12, carries redundant information as IntraMbMode. The notation I_16x16_x_y_z used in the table, 'x' is Intra16x16LumaPredMode, 'y' is ChromaCbpInd, and 'z' is LumaCbpInd, as defined in Table 3-13.



Table 3-12. MbType definition for Intra Macroblock

| Macroblock Type | MbType |
|-----------------|------------------|
| I_4x4 | 0 |
| I_8x8 | 0 |
| I_16x16_0_0_0 | 1 |
| I_16x16_1_0_0 | 2 |
| I_16x16_2_0_0 | 3 |
| I_16x16_3_0_0 | 4 |
| I_16x16_0_1_0 | 5 |
| I_16x16_1_1_0 | 6 |
| I_16x16_2_1_0 | 7 |
| I_16x16_3_1_0 | 8 |
| I_16x16_0_2_0 | 9 |
| I_16x16_1_2_0 | Ah |
| I_16x16_2_2_0 | Bh |
| I_16x16_3_2_0 | Ch |
| I_16x16_0_0_1 | Dh |
| I_16x16_1_0_1 | Eh |
| I_16x16_2_0_1 | Fh |
| I_16x16_3_0_1 | 10h |
| I_16x16_0_1_1 | 11h |
| I_16x16_1_1_1 | 12h |
| I_16x16_2_1_1 | 13h |
| I_16x16_3_1_1 | 14h |
| I_16x16_0_2_1 | 15h |
| I_16x16_1_2_1 | 16h |
| I_16x16_2_2_1 | 17h |
| I_16x16_3_2_1 | 18h |
| I_PCM | 19h (used by HW) |

For Intra_16x16 modes, the 5 bits of value (MbType – 1) have the following meanings.



Table 3-13. Sub field definition used by MbType for a macroblock with Intra16x16 prediction

| Bits | Description |
|------|--|
| 4 | <p>LumaCbplnd – Luma Coded Block Pattern Indicator 0 means none of the luma blocks are coded. 1 means that at least one luma block is coded. 0 = SUBMODE_I16_L_0 1 = SUBMODE_I16_L_NZ In VME output, this field is forced to be 1 before adding 1 in Intra_16x16 mode.</p> |
| 3:2 | <p>ChromaCbplnd – Chroma Coded Block Pattern Indicator 00 means none of chroma blocks are coded. 01 means that only the chroma DC block is coded, but all AC blocks are not coded. 10 means that at least one AC chroma block is coded. 00 = SUBMODE_I16_C_0 01 = SUBMODE_I16_C_DC 10 = SUBMODE_I16_C_NZ 11 = Reserved In VME output, this field is forced to be 10 before adding 1 in Intra_16x16 mode. <i>Programming Note: Adding 1 to MbType by VME hardware may have carry in to this field. But as '11' is reserved, the carry-in doesn't propagate into bit 4 or higher. This allows software to update MbType, if desired, using the redundant LumaIntraPredModes information.</i></p> |
| 1:0 | <p>Intra16x16PredMode – Intra16x16 Prediction Mode These two bits carries redundant (identical) information as that in LumaIntraPredModes[0][0]. 0 = SUBMODE_I16_VER 1 = SUBMODE_I16_HOR 2 = SUBMODE_I16_DC 3 = SUBMODE_I16_PLANE</p> |

Table 3-14. IntraMbMode definition

| IntraMbMode [1:0] | Description | Supported by VME? | Used by PAK? |
|-------------------|--|-------------------|--------------|
| 0 | INTRA_16x16 (redundant with MbType) | Yes | Ignored |
| 1 | INTRA_8x8 | Yes | Yes |
| 2 | INTRA_4x4 | Yes | Yes |
| 3 | IPCM (redundant with MbType) | No | Ignored |



As an alternative representation, MbType is logically the same as the following, except the I_PCM and I_NxN (i.e. I_4x4 and I_8x8) cases:

- 24 types of 16x16 intra modes: **A+B+C+D:** (1h – 18h)

| | | |
|--------------------|----|---|
| MBTYPE_INTRA_16x16 | 1h | A |
|--------------------|----|---|

- 4 Intra16x16 modes:

| | | |
|-----------------|---|---|
| SUBMODE_I16_VER | 0 | B |
| SUBMODE_I16_HOR | 1 | B |
| SUBMODE_I16_DC | 2 | B |
| SUBMODE_I16_PLN | 3 | B |

- 3 Chroma Cbp indices:

| | | |
|------------------|---|---|
| SUBMODE_I16_C_0 | 0 | C |
| SUBMODE_I16_C_DC | 4 | C |
| SUBMODE_I16_C_NZ | 8 | C |

- 2 Luma Cbp indices:

| | | |
|------------------|----|---|
| SUBMODE_I16_L_0 | 0 | D |
| SUBMODE_I16_L_NZ | Ch | D |

3.6.4.1.5 Macroblock Type for Inter Cases

Sub-Macroblock Prediction Mode, SubMbPredMode, indicates the prediction mode for the sub-partitions. Prediction mode specifies prediction direction being forward (from L0), backward (from L1) or bi-directional (from both L0 and L1). Its meaning depends on InterMbMode. Table 3-17 provides the definition of the field.

- If InterMbMode is INTER16x16, only SubMbPredMode[0] is valid, it describes the prediction mode of the 16x16 macroblock. The other entries are set to zero by hardware.
 - For AVC, SubMbPredMode[0] contains redundant information as encoded in MbType parameter.
 - *Note: SubMbPredMode[1]-[3] are intentionally set to zero to allow a simple LUT to derive MbType as described later.*
- If InterMbMode is INTER16x8, and INTER8x16, only the first two entries SubMbPredMode[0] and SubMbPredMode[1] are valid, describing the sub-macroblock prediction mode.
 - For AVC, SubMbPredMode[0]/[1] contains redundant information as encoded in MbType parameter.
 - *Note: SubMbPredMode[2]-[3] are intentionally set to zero to allow a simple LUT to derive MbType as described later.*
- If InterMbMode is INTER8x8, each entry of SubMbPredMode describes the prediction mode of the sub-partition of an 8x8 sub-macroblock.
 - For AVC, SubMbPredMode can be derived from sub_mb_type field for BP_8x8 macroblocks as defined in AVC spec.



- Note on Direct Sub-macroblock Prediction Mode: Direct prediction is not conveyed through SubMbPredMode, instead, it is carried through Direct8x8Pattern.

Table 3-15. InterMbMode definition

| MbSkipFlag | InterMbMode | Description |
|--------------------|-----------------------|--------------------------|
| 0 | 0 | INTER16x16 |
| 0 | 1 | INTER16x8 |
| 0 | 2 | INTER8x16 |
| 0 | 3 | INTER8x8 |
| 1 | 0 | PSKIP/BSKIP16x16* |
| 1 | 3 | BSKIP |
| 1 | 1, 2 | Reserved |
| Used by PAK | Ignored by PAK | |

* BSKIP16x16 is an optional non-standard but equivalent optimization.

Table 3-16. Definition of SubMbPredMode based on InterMbMode

| SubMbPredMode | INTER16x16 | INTER16x8 | INTER8x16 | INTER8x8 |
|---------------|-----------------------|-----------------------|-----------------------|--------------------|
| Bit | MbType = [1...3] | MbType = [16h] | MbType = [4...15h] | MbType = [16h] |
| 7:6 | MBZ | MBZ | MBZ | Block8x8 3 |
| 5:4 | MBZ | MBZ | MBZ | Block8x8 2 |
| 3:2 | MBZ | Block16x8 1 | Block8x16 1 | Block8x8 1 |
| 1:0 | Block16x16 | Block16x8 0 | Block8x16 0 | Block8x8 0 |
| | Ignored by PAK | Ignored by PAK | Ignored by PAK | Used by PAK |

Table 3-17. Definition of SubMbPredMode[i]

| SubMbPredMode | Description | InterMbMode | VME Output | MvCountPred | Comments |
|---------------|-------------|-------------|------------|-------------|--------------|
| 0 | Pred_L0 | All | Yes | 1 | P or B Slice |
| 1 | Pred_L1 | All | Yes | 1 | B Slice Only |
| 2 | BiPred | All | Yes | 2 | B Slice Only |
| 3 | Reserved | Reserved | Reserved | Reserved | Reserved |

Sub-Macroblock Shape, SubMbShape[i], for $i = 0 \dots 3$, describes the shape of the sub partitions of the 8x8 sub-macroblock of a BP_8x8 macroblock. This field is only valid if InterMBMode is INTER8x8. They are defined in Table 3-18. The parameters can be derived from *sub_mb_type* field as defined in AVC spec.



Note: These fields must be correctly set even for **Direct** or **Skip** 8x8 cases, the individual B_Direct_8x8 block is flagged by the **Direct8x8Pattern** variable.

Table 3-18. Definition of SubMbShape for an 8x8 region of a BP_8x8 macroblock (including BSKIP, BDIRECT)

| SubMbShape | Description | | | |
|------------|--------------|----------------|-----------------|--------------|
| | NumSubMbPart | SubMbPartWidth | SubMbPartHeight | MvCountShape |
| 0 | 1 | 8 | 8 | 1 |
| 1 | 2 | 8 | 4 | 2 |
| 2 | 2 | 4 | 8 | 2 |
| 3 | 4 | 4 | 4 | 4 |

For an inter macroblock, MbType, carries redundant information as InterMbMode and SubMbPredMode. Table 3-19 provides the typical inter macroblock types and Table 3-20 provides that with skip and direct modes. The definition of MbType for both P slice and B slice is the same and is equivalent to that for mb_type of a B slice in the AVC spec. As direct mode is indicated using a separate field Direct8x8Pattern, 0 is reserved for MbType.

Here, MVCount is the number of motion vectors actually encoded in the bitstream. It is informative. For a BP_8x8 or equivalent Skip/Direct macroblock, MVCount is the sum of the following term for the four 8x8 sub macroblock (with $i = 0 \dots 3$):

$$MvCountShape[i] * MvCountPred[i] * MvCountDirect[i]$$

where MvCountShape[i] is block count for sub macroblock [i], MvCountPred[i] is the motion vector count for each block of sub macroblock[i], and MvCountDirect[i] is the multiplier for direct mode for B Slice, indicating whether motion vectors are coded or not. It must be set to 1 for P slice. For B Slice, $MvCountDirect[i] = !Direct8x8Pattern[i]$, which is 0 for a sub macroblock coded as direct mode and 1 otherwise.

In the tables, “DC” stands for “Don’t Care” as PAK hardware ignores these fields.

Table 3-19. MbType definition for Inter Macroblock (and MbSkipflag = 0)

| Macroblock Type | MbType | MbSkipFlag | Direct8x8Pattern | SubMbShape | SubMbPredMode | MVCount |
|----------------------|--------|------------|------------------|------------|---------------|---------|
| <i>Reserved</i> | 0 | - | - | - | - | - |
| <i>BP_LO_16x16</i> | 1 | 0 | 0 | DC | DC | 1 |
| <i>B_L1_16x16</i> | 2 | 0 | 0 | DC | DC | 1 |
| <i>B_Bi_16x16</i> | 3 | 0 | 0 | DC | DC | 2 |
| <i>BP_LO_LO_16x8</i> | 4 | 0 | 0 | DC | DC | 2 |
| <i>BP_LO_LO_8x16</i> | 5 | 0 | 0 | DC | DC | 2 |
| <i>B_L1_L1_16x8</i> | 6 | 0 | 0 | DC | DC | 2 |
| <i>B_L1_L1_8x16</i> | 7 | 0 | 0 | DC | DC | 2 |
| <i>B_LO_L1_16x8</i> | 8 | 0 | 0 | DC | DC | 2 |
| <i>B_LO_L1_8x16</i> | 9 | 0 | 0 | DC | DC | 2 |



| Macroblock Type | MbType | MbSkipFlag | Direct8x8Pattern | SubMbShape | SubMbPredMode | MVCount |
|-----------------|---------|------------|------------------|------------|---------------|---------|
| B_L1_L0_16x8 | 0Ah | 0 | 0 | DC | DC | 2 |
| B_L1_L0_8x16 | 0Bh | 0 | 0 | DC | DC | 2 |
| B_L0_Bi_16x8 | 0Ch | 0 | 0 | DC | DC | 3 |
| B_L0_Bi_8x16 | 0Dh | 0 | 0 | DC | DC | 3 |
| B_L1_Bi_16x8 | 0Eh | 0 | 0 | DC | DC | 3 |
| B_L1_Bi_8x16 | 0Fh | 0 | 0 | DC | DC | 3 |
| B_Bi_L0_16x8 | 10h | 0 | 0 | DC | DC | 3 |
| B_Bi_L0_8x16 | 11h | 0 | 0 | DC | DC | 3 |
| B_Bi_L1_16x8 | 12h | 0 | 0 | DC | DC | 3 |
| B_Bi_L1_8x16 | 13h | 0 | 0 | DC | DC | 3 |
| B_Bi_Bi_16x8 | 14h | 0 | 0 | DC | DC | 4 |
| B_Bi_Bi_8x16 | 15h | 0 | 0 | DC | DC | 4 |
| BP_8x8 | 16h | 0 | != F | vary | vary | Sum |
| <i>Reserved</i> | 17h-1Fh | - | - | - | - | - |

Table 3-20. Additional MbType definition with Direct/Skip for Inter Macroblock

| Macroblock Type | MbType | Xfrm 8x8 | MbSkipFlag | Direct8x8Pattern | SubMbShape | SubMbPredMode | MvCount | Comments |
|----------------------------|--------|----------|------------|------------------|------------|---------------|---------|--|
| P_Skip_16x16 | 1 | - | 1 | DC | DC | DC | 0 | Skipped macroblock. Motion compensation like P_L0_16x16 |
| B_Skip_16x16_4MVP | 16h | Vary | 1 | Fh | 0 | vary | 0 | Skipped macroblock. Motion compensation like B_8x8 with 8x8 subblocks, when direct_8x8_inference_flag is set to 1 |
| B_Skip_16x16_16MVP | 16h | 0 | 1 | Fh | FFh | vary | 0 | Skipped macroblock. Motion compensation like B_8x8 with 4x4 subblocks, when direct_8x8_inference_flag is set to 0 |
| B_Direct_16x16_4MVP | 16h | vary | 0 | Fh | 0 | vary | 0 | MbType coded as B_Direct_16x16. Motion compensation like B_8x8 with 8x8 subblocks, when direct_8x8_inference_flag is set to 1 |



| | | | | | | | | |
|------------------------------|-----|---|---|----|-----|------|---|--|
| B_Direct_16x16_16 MVP | 16h | 0 | 0 | Fh | FFh | vary | 0 | MbType coded as B_Direct_16x16. Motion compensation like B_8x8 with 4x4 subblocks, when direct_8x8_inference_flag is set to 0 |
|------------------------------|-----|---|---|----|-----|------|---|--|

People might notice that B_DIRECT_16x16 and B_SKIP are mapped on BP_8x8 for AVC decoding interface in IT mode as the motion compensation operation for both modes are the same as BP_8x8. According to AVC Spec, motion vectors for B_DIRECT_16x16 and B_SKIP are derived from temporally co-located macroblock on an 8x8 sub macroblock basis if `direct_8x8_inference_flag` is set to 1 or on a 4x4 block basis if it is set to 0. For each sub macroblock or block, `SubMbPredMode` is derived, thus can any of the valid numbers. Motion vectors may also be different. In spatial direct mode, the motion vectors are subject to spatial neighbor macroblocks as well as co-located macroblock. The spatial prediction is based on the neighbor macroblocks, so the same spatial predicted motion vector applies to all sub macroblocks or blocks. However, under certain conditions, temporal predictor may replace (`colZeroFlag`) the spatial predictor for a given sub macroblock or block. Thus the motion vectors may differ.

In Table 3-19, the macroblock type names for major partitions nicely follow forms *BP_MbPredMode_MbShape* (like BP_L0_16x16) and *B_MbPredMode0_MbPredMode1_MbShape* (like B_L0_Bi_16x8). For minor partitions it is fixed as *BP_MbShape* as BP_8x8.



However, in Table 3-20 the macroblock types for Skip and Direct modes does not follow the same rule. The third field in P_Skip_16x16 or B_Direct_16x16_x indicates that “Skip” or “Direct” applies to the entire 16x16 macroblock, even though MbShape is 8x8 as that in BP_8x8. In order to distinguish the SubMbShape being 8x8 or 4x4 for B_Skip and B_Direct, the fourth field is added. 4MVP indicates upto 4 MV pairs are presented with SubMbShape equals to 0; and 16MVP indicates up to 16 MV pairs are presented with SubMbShape equals to FFh. Also note that P_8x8ref0 is not specified in PAK input interface, it is up to hardware to detect and choose its packing format based on number of reference indices and reference index for the given macroblock.

3.6.4.1.6 Macroblock Type Conversion Rules

For improved coding efficiency the PAK hardware has the capability to convert macroblock types to use more efficiency coding modes such as DIRECT and SKIP. For an inter macroblock or a sub macroblock coded as DIRECT, no motion vector is needed in the bitstream for the macroblock or sub macroblock. If a macroblock is coded as SKIP, it only consumes one SKIP bit (no motion vector, no coefficients are coded). And information about the macroblock is ‘inferred’ according to the rules stated in the AVC Spec.

As the input to PAK, the following signals can convey the information regarding DIRECT and SKIP:

- MbSkipFlag
- Direct8x8Pattern
- CodecBlockPattern (CbpY, CbpCb, CbpCr)

Such conversion can be enabled or disabled through the SLICE_STATE fields **DirectConvDisable** and **SkipConvDisable** as well as the in line command field **MbSkipConvDisable**.

A P slice doesn’t support direct mode, it only supports P_Skip, which is equivalent to a 16_16_L0 prediction. Other prediction types cannot be converted to P_Skip. The following table describes the macroblock type conversion rules for a P slice. Here CBP = CbpY/CbpCb/CbpCr are the final computed results after quantization by the hardware. Note that hardware honors the input CbpY/CbpCb/CbpCr fields – if the value corresponding to a block is set to zero, the resulting CBP is also zero. The output **mb_skip_flag** and **mb_type** are the symbols coded in the bitstream as defined in the AVC spec. “DC” stands for “Don’t care”, “T” for “True”.

Note that the internal condition of MV==MVP is subject to the precise rules stated in the AVC Spec as quoted below. Note that there are exceptions for P_Skip from the normal motion vector prediction rules.

Derivation process for luma motion vectors for skipped macroblocks in P and SP slices

This process is invoked when mb_type is equal to P_Skip.

Outputs of this process are the motion vector mvL0 and the reference index refIdxL0.

The reference index refIdxL0 for a skipped macroblock is derived as follows.

$$refIdxL0 = 0. \tag{8-168}$$

For the derivation of the motion vector mvL0 of a P_Skip macroblock type, the following applies.

– *The process specified in subclause 8.4.1.3.2 is invoked with mbPartIdx set equal to 0, subMbPartIdx set equal to 0, currSubMbType set equal to "na", and listSuffixFlag set equal to 0 as input and the output is assigned to mbAddrA, mbAddrB, mvLOA, mvLOB, refIdxLOA, and refIdxLOB.*



- The variable *mvL0* is specified as follows.
- If any of the following conditions are true, both components of the motion vector *mvL0* are set equal to 0.
 - *mbAddrA* is not available
 - *mbAddrB* is not available
 - *refIdxLOA* is equal to 0 and both components of *mvLOA* are equal to 0
 - *refIdxLOB* is equal to 0 and both components of *mvLOB* are equal to 0
 - Otherwise, the derivation process for luma motion vector prediction as specified in subclause 8.4.1.3 is invoked with *mbPartIdx* = 0, *subMbPartIdx* = 0, *refIdxL0*, and *currSubMbType* = "na" as inputs and the output is assigned to *mvL0*.

NOTE – The output is directly assigned to *mvL0*, since the predictor is equal to the actual motion vector.

Table 3-21. Macroblock type conversion rule for an inter macroblock in a P slice

| Input | | Internal | | | Output | | Notes |
|-----------------|------------------------------------|----------|-----------|------------------|--------------|---------|--|
| Macroblock Type | SkipConvDisable SkipConvDisable | CBP | MV == MVP | MbAffSkipAllowed | mb_skip_flag | mb_type | |
| P_Skip_16x16 | DC | DC | DC | 1 | 1 | - | Forced to P_Skip ; Hardware will force CBP to zero and also ignore SkipConvDisable control. Hardware doesn't check for MV==MVP error condition |
| P_Skip_16x16 | DC | DC | DC | 0 | 0 | 0 | Reverse convert to P_L0_16x16 ; Hardware will force CBP to zero but reversely convert MbType as P_L0_16x16 once it determines that Skip is not allowed. |
| BP_16x16_L0 | 0 | 0 | T | 1 | 1 | - | Converted to P_Skip . Even input doesn't provide skip hint, hardware can performance the optimization by detecting CBP and MV==MVP condition. |
| BP_16x16_L0 | 0 | 0 | T | 0 | 0 | 0 | Reverse back to P_L0_16x16 ; Hardware will reverse back to P_L0_16x16 even Skip conditions are met once it determines that Skip is not allowed. |
| BP_16x16_L0 | 1 | 0 | T | T | 0 | 0 | Still coded as P_L0_16x16 = 0. |

A B slice supports both direct and skip modes. The following table describes the macroblock type conversion rules for a B slice. Hardware does not verify MV==MVP condition for a Skip/Direct macroblock in a B Slice as no DMV is performed by hardware.



Table 3-22. Macroblock type conversion rule for an inter macroblock in a B slice

| Input | | | Internal | | | Output | | Notes |
|-------------------------------------|------------------------------------|-------------------|----------|-----------|------------------|--------------|---------|---|
| Macroblock Type | SkipConvDisable SkipConvDisable | DirectConvDisable | CBP | MV == MVP | MbAffSkipAllowed | mb_skip_flag | mb_type | |
| B_Skip_8x8 B_Skip_4x4 | DC | DC | DC | n/a | 1 | 1 | - | Forced to B_Skip; Hardware will force CBP to zero and also ignore SkipConvDisable control. |
| B_Skip_8x8 B_Skip_4x4 | DC | DC | DC | n/a | 0 | 0 | 0 | REVERSE convert to B_Direct_16x16; Hardware will force CBP to zero and also reverse convert to B_Direct_16x16 when it discovers Skip is not allowed. |
| B_Direct_16x16_4 MVPair/16MVPair | 0 | 0 | 0 | n/a | 1 | 1 | - | Converted to B_Skip. Hardware first converts to B_Direct_16x16 and then further to B_Skip if CBP = 0. |
| B_Direct_16x16_4 MVPair/16MVPair | 0 | 0 | 0 | n/a | 0 | 0 | 0 | Converted to B_Direct_16x16. Hardware first converts to B_Direct_16x16 and stop there as it discovers Skip is not allowed even CBP=0. |
| B_Direct_16x16_4 MVPair/16MVPair | 1 | 0 | 0 | n/a | DC | 0 | 0 | Converted to B_Direct_16x16. Hardware converts to B_Direct_16x16 and stops there even though CBP = 0 as input disallows Skip conversion. |
| B_Direct_16x16_4 MVPair/16MVPair | DC | 0 | NZ | n/a | DC | 0 | 0 | Converted to B_Direct_16x16. Hardware converts to B_Direct_16x16 and stops there because CBP != 0. |



| Input | | | Internal | | | Output | | Notes |
|-------------------------------------|------------------------------------|-------------------|----------|-----------|------------------|--------------|---------|---|
| Macroblock Type | SkipConvDisable SkipConvDisable | DirectConvDisable | CBP | MV == MVP | MbAffSkipAllowed | mb_skip_flag | mb_type | |
| B_Direct_16x16_4 MVPair/16MVPair | DC | 1 | DC | n/a | DC | 0 | 16h | Stay as B_8x8. Hardware stays at B_8x8 and codes each sub macroblocks even all are direct. |



The internal signal **MbAffSkipAllowed** is added to deal with a restriction on the frame/field flag (**MbFieldFlag**) which is unique to MBAFF. **MbAffSkipAllowed** is always set to 1 in non-MBAFF modes. In MBAFF mode, a macroblock pair may be both skipped only if its **MbFieldFlag** is the same as its available neighbor macroblock pair A or B if A or B is available (in that order), or is not 0 if A/B are both not available. Otherwise, one of the macroblocks in the pair must be coded.

To reduce the burden on software, PAK hardware handles the above restriction correctly. For the first MB in a pair, **MbAffSkipAllowed** is always set to 1. Therefore, hardware allows converting the first MB to Skip if skip conversion is enabled. For the second MB in a pair, hardware sets **MbAffSkipAllowed** to 0 if the following is true:

- The current MB Pair has different **MbFieldFlag** than its available neighbor A or B if A or B is available, or is not 0 if A/B are both not available
- And the first MB is coded as a SKIP (could be forced or converted)

Otherwise, it sets **MbAffSkipAllowed** to 1. As **MbAffSkipAllowed** is to 0 for the above condition, hardware will disallow Skip mode for the second MB. If the input signal forces it to Skip, hardware performs reverse-conversion to code it as P_L0_16x16 or B_Direct_16x16 with CBP = 0 for a macroblock in a P or B Slice. This means that hardware is able to correct the programming mistake by software. If the macroblock is not forced to skip, hardware simply disallows Skip conversion.

Software still has an option to disallow Skip Conversion on a per-MB basis using the **MbSkipConvDisable** control field in the inline command.

3.6.4.2 Indirect Data Description

For each macroblock, an ENC-PAK data set consists of two types of data blocks: indirect **MV data block** and **inline MB information**.

The indirect MV data block may be in two modes: **unpacked mode** and **packed-size mode**.

3.6.4.2.1 Unpacked Motion Vector Data Block

In the **unpacked** mode, motion vectors are expanded (or duplicated) to either bidirectional 8x8 8MV major partition format, or bidirectional 4x4 32MV format. Thus either 32 bytes or 128 bytes is assigned to each MB.

Motion Vector block contains motion vectors in an intermediate format that is partially expanded according to the sub-macroblock size. During the expansion, a place that does not contain a motion vector is filled by replicating the relevant motion vector according to the following motion vector replication rules. If the relevant motion vector doesn't exist (for the given L0 or L1), it is zero filled.



Motion Vector Replication Rules:

- Rule #1
 - #1.1: For L0 MV, for any sub-macroblock or sub-partition where there is at least one motion vector
 - If L0 inter prediction exists, the corresponding L0 MV is used
 - Else it must be zero
 - #1.2: For L1 MV, for any sub-macroblock or sub-partition where there is at least one motion vector
 - If L1 inter prediction exists, the corresponding L1 MV is used
 - Else it must be zero
- For a macroblock with a 16x16, 16x8 or 8x16 sub-macroblock, MvSize = 8. The eight MV fields follow Rule #1.
 - The 16x16 is broken down into 4 8x8 sub-macroblocks. The 16x16 MVs (after rule #1) are replicated into all 8x8 blocks.
 - For an 8x16 partition, each 8x16 is broken down into 2 8x8 stacking vertically. The 8x16 MVs (after rule #1) are replicated into both 8x8 blocks.
 - For a 16x8 partition, each 16x8 is broken down into 2 8x8 stacking horizontally. The 16x8 MVs (after rule #1) are replicated into both 8x8 blocks.
- For macroblock with sub-macroblock of 8x8 without minor partition (SubMbShape[0...3] = 0), MvSize = 8, (e.g. mb_type equal to P_8x8, P_8x8ref0, or B_8x8)
 - There is no motion vector replication
- For macroblock with sub-macroblock of 8x8 with at least one minor partition (if any SubMbShape[i] != 0), MvSize = 32, (e.g. mb_type equal to P_8x8, P_8x8ref0, or B_8x8)
 - For an 8x8 sub-partition, the 8x8 MVs (after rule #1) is replicated into all the four 4x4 blocks.
 - For an 4x8 sub-partition within an 8x8 partition, each 4x8 is broken down into 2 4x4 stacking vertically. The 4x8 MVs (after rule #1) are replicated into both 4x4 blocks.
 - For an 8x4 sub-partition within an 8x8 partition, each 8x4 is broken down into 2 4x4 stacking horizontally. The 8x4 MVs (after rule #1) are replicated into both 4x4 blocks.
 - For a 4x4 sub-partition within an 8x8 partition, each 4x4 has its own MVs (after rule #1).



Table 3-23. Motion Vector block and MvSize

| DWord | Bit | MvSize | |
|-------|-------|------------|--------------|
| | | 8 | 32 |
| W1.0 | 31:16 | MV_Y0_L0.y | MV_Y0_0_L0.y |
| | 15:0 | MV_Y0_L0.x | MV_Y0_0_L0.x |
| W1.1 | 31:16 | MV_Y0_L1.y | MV_Y0_0_L1.y |
| | 15:0 | MV_Y0_L1.x | MV_Y0_0_L1.x |
| W1.2 | 31:0 | MV_Y1_L0 | MV_Y0_1_L0 |
| W1.3 | 31:0 | MV_Y1_L1 | MV_Y0_1_L1 |
| W1.4 | 31:0 | MV_Y2_L0 | MV_Y0_2_L1 |
| W1.5 | 31:0 | MV_Y2_L1 | MV_Y0_2_L0 |
| W1.6 | 31:0 | MV_Y3_L0 | MV_Y0_3_L0 |
| W1.7 | 31:0 | MV_Y3_L1 | MV_Y0_3_L1 |
| W2.0 | 31:0 | n/a | MV_Y1_0_L1 |
| W2.1 | 31:0 | n/a | MV_Y1_0_L0 |
| W2.2 | 31:0 | n/a | MV_Y1_1_L1 |
| W2.3 | 31:0 | n/a | MV_Y1_1_L0 |
| W2.4 | 31:0 | n/a | MV_Y1_2_L1 |
| W2.5 | 31:0 | n/a | MV_Y1_2_L0 |
| W2.6 | 31:0 | n/a | MV_Y1_3_L0 |
| W2.7 | 31:0 | n/a | MV_Y1_3_L1 |
| W3.0 | 31:0 | n/a | MV_Y2_0_L1 |
| W3.1 | 31:0 | n/a | MV_Y2_0_L0 |
| W3.2 | 31:0 | n/a | MV_Y2_1_L1 |
| W3.3 | 31:0 | n/a | MV_Y2_1_L0 |
| W3.4 | 31:0 | n/a | MV_Y2_2_L1 |
| W3.5 | 31:0 | n/a | MV_Y2_2_L0 |
| W3.6 | 31:0 | n/a | MV_Y2_3_L0 |
| W3.7 | 31:0 | n/a | MV_Y2_3_L1 |



| DWord | Bit | MvSize | |
|-------|------|--------|------------|
| | | 8 | 32 |
| W4.0 | 31:0 | n/a | MV_Y3_0_L1 |
| W4.1 | 31:0 | n/a | MV_Y3_0_L0 |
| W4.2 | 31:0 | n/a | MV_Y3_1_L1 |
| W4.3 | 31:0 | n/a | MV_Y3_1_L0 |
| W4.4 | 31:0 | n/a | MV_Y3_2_L1 |
| W4.5 | 31:0 | n/a | MV_Y3_2_L0 |
| W4.6 | 31:0 | n/a | MV_Y3_3_L0 |
| W4.7 | 31:0 | n/a | MV_Y3_3_L1 |

The motion vector(s) for a given sub-macroblock or a sub-partition are uniquely placed in the output message as shown by the non-duplicate fields in Table 3-24 and Table 3-25.

MV_Yx_L0 and MV_Yx_L1 may be present individually or both. If one is not present, the corresponding field must be zero. Subsequently, the duplicated fields will be zero as well.

Table 3-24. Motion Vector duplication by sub-macroblocks for a 16x16 macroblock, whereas the 8x8 column is for 4x(8x8) partition without minor shape

| DWord | Bit | 16x16 | 16x8 | 8x16 | 8x8 |
|-------|-------|---------------|---------------|---------------|----------|
| | | | | | |
| W1.0 | 31:16 | MV_Y0_L1 (A) | MV_Y0_L1 (A) | MV_Y0_L1 | MV_Y0_L1 |
| | 15:0 | MV_Y0_L0 (A) | MV_Y0_L0 (A) | MV_Y0_L0 | MV_Y0_L0 |
| W1.1 | 31:16 | Duplicate (A) | Duplicate (A) | MV_Y1_L1 | MV_Y1_L1 |
| | 15:0 | Duplicate (A) | Duplicate (A) | MV_Y1_L0 | MV_Y1_L0 |
| W1.2 | 31:16 | Duplicate (A) | MV_Y2_L1 (B) | Duplicate (A) | MV_Y2_L1 |
| | 15:0 | Duplicate (A) | MV_Y2_L0 (B) | Duplicate (A) | MV_Y2_L0 |
| W1.3 | 31:16 | Duplicate (A) | Duplicate (B) | Duplicate (B) | MV_Y3_L1 |
| | 15:0 | Duplicate (A) | Duplicate (B) | Duplicate (B) | MV_Y3_L0 |



Table 3-25. Motion Vector duplication by sub-partitions for the first 8x8 sub-macroblock Y0 if any Y0-Y3 contains minor shape (Y1_ to Y3_ have the same format in W2 to W4)

| DWord | Bit | | | | |
|-------|-------|---------------|----------------|----------------|------------|
| | | 8x8 | 8x4 | 4x8 | 4x4 |
| W1.0 | 31:16 | MV_Y0_L1 | MV_Y0_0_L1 (A) | MV_Y0_0_L1 (A) | MV_Y0_0_L1 |
| | 15:0 | MV_Y0_L0 | MV_Y0_0_L0 (A) | MV_Y0_0_L0 (A) | MV_Y0_0_L0 |
| W1.1 | 31:16 | Duplicate (A) | Duplicate (A) | MV_Y0_1_L1 (B) | MV_Y0_1_L1 |
| | 15:0 | Duplicate (A) | Duplicate (A) | MV_Y0_1_L0 (B) | MV_Y0_1_L0 |
| W1.2 | 31:16 | Duplicate (A) | MV_Y0_2_L1 (B) | Duplicate (A) | MV_Y0_2_L1 |
| | 15:0 | Duplicate (A) | MV_Y0_2_L0 (B) | Duplicate (A) | MV_Y0_2_L0 |
| W1.3 | 31:16 | Duplicate (A) | Duplicate (B) | Duplicate (B) | MV_Y0_3_L0 |
| | 15:0 | Duplicate (A) | Duplicate (B) | Duplicate (B) | MV_Y0_3_L1 |

3.6.4.2.2 Packed-size Motion Vector Data Block

In the packed case, no redundant motion vectors are sent. So the number of motion vectors sent, as specified by **MvQuantity** is the same as the motion vectors that will be packed (**MvPacked**).

The following tables are for information only. Fields like MvQuantity and MvPacked are not required interface fields.

| MbSkipFlag | MbType | Description | Mv Quantity | MvSize | (Minimal MvSize) |
|------------|--------|---------------|-------------|--------|------------------|
| 1 | 1 | P_Skip_16x16 | 0 | 8 | 1 |
| 0 | 1 | BP_L0_16x16 | 1 | 8 | 1 |
| 0 | 2 | B_L1_16x16 | 1 | 8 | 1 |
| 0 | 3 | B_Bi_16x16 | 2 | 8 | 2 |
| 0 | 4 | BP_L0_L0_16x8 | 2 | 8 | 4 |
| 0 | 5 | BP_L0_L0_8x16 | 2 | 8 | 4 |
| 0 | 6 | B_L1_L1_16x8 | 2 | 8 | 8 |
| 0 | 7 | B_L1_L1_8x16 | 2 | 8 | 8 |
| 0 | 8 | B_L0_L1_16x8 | 2 | 8 | 8 |
| 0 | 9 | B_L0_L1_8x16 | 2 | 8 | 8 |
| 0 | 0Ah | B_L1_L0_16x8 | 2 | 8 | 8 |
| 0 | 0Bh | B_L1_L0_8x16 | 2 | 8 | 8 |
| 0 | 0Ch | B_L0_Bi_16x8 | 3 | 8 | 8 |
| 0 | 0Dh | B_L0_Bi_8x16 | 3 | 8 | 8 |



| | | | | | |
|---|-----|---------------|----|---------|---------|
| 0 | 0Eh | B_L1_Bi_16x8 | 3 | 8 | 8 |
| 0 | 0Fh | B_L1_Bi_8x16 | 3 | 8 | 8 |
| 0 | 10h | B_Bi_L0_16x8 | 3 | 8 | 8 |
| 0 | 11h | B_Bi_L0_8x16 | 3 | 8 | 8 |
| 0 | 12h | B_Bi_L1_16x8 | 3 | 8 | 8 |
| 0 | 13h | B_Bi_L1_8x16 | 3 | 8 | 8 |
| 0 | 14h | B_Bi_Bi_16x8 | 4 | 8 | 8 |
| 0 | 15h | B_Bi_Bi_8x16 | 4 | 8 | 8 |
| 0 | 16h | BP_8x8 | ≥4 | 8 or 32 | 8 or 32 |

When MbType = 22, BP_8x8, take the sum of four individual 8x8 subblocks

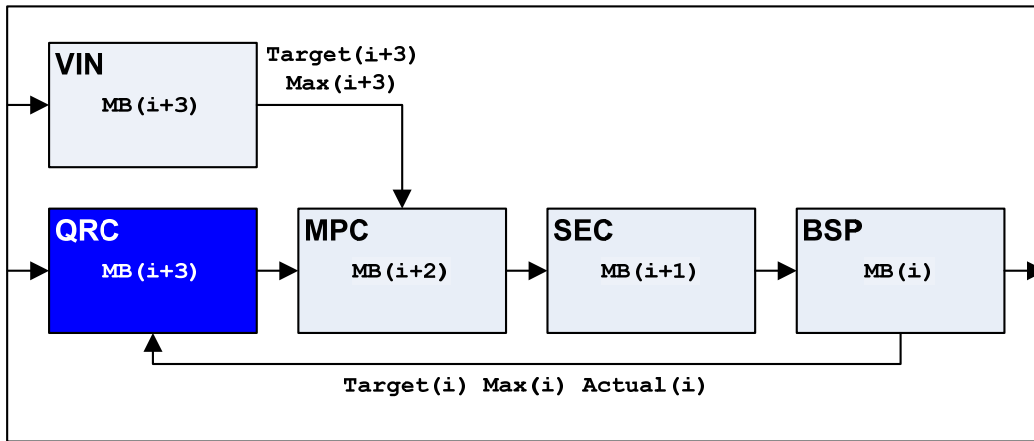
| Direct8x8 Pattern | SubMb Shape | SubMb PredMode | Description | Mv Quantity | Mv Size | (Min MvSize) |
|-------------------|-------------|----------------|--|-------------|---------|--------------|
| OR | OR | OR | | ADD | ADD | ADD |
| 1 | 0 | 0 | P_Skip_8x8 B_Direct_L0_8x8 (B-Skip_L0_8x8) | 0 | 2 | 1 |
| 1 | 0 | 1 | B_Direct_L1_8x8 (B-Skip_L1_8x8) | 0 | 2 | 1 |
| 1 | 0 | 2 | B_Direct_Bi_8x8 (B-Skip_Bi_8x8) | 0 | 2 | 2 |
| 1 | 3 | 0 | P_Skip_4x4 B_Direct_L0_4x4 (B-Skip_L0_4x4) | 0 | 8 | 4 |
| 1 | 3 | 1 | B_Direct_L1_4x4 (B-Skip_L1_4x4) | 0 | 8 | 4 |
| 1 | 3 | 2 | B_Direct_Bi_4x4 (B-Skip_Bi_4x4) | 0 | 8 | 8 |
| 0 | 0 | 0 | BP_L0_8x8 | 1 | 2 | 1 |
| 0 | 0 | 1 | B_L1_8x8 | 1 | 2 | 1 |
| 0 | 0 | 2 | B_Bi_8x8 | 2 | 2 | 2 |
| 0 | 1 | 0 | BP_L0_8x4 | 2 | 8 | 4 |
| 0 | 1 | 1 | B_L1_8x4 | 2 | 8 | 4 |
| 0 | 1 | 2 | B_Bi_8x4 | 4 | 8 | 8 |
| 0 | 2 | 0 | BP_L0_4x8 | 2 | 8 | 4 |
| 0 | 2 | 1 | B_L1_4x8 | 2 | 8 | 4 |



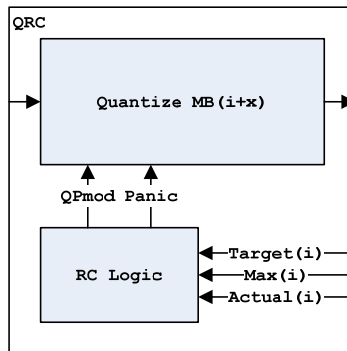
| Direct8x8 Pattern | SubMb Shape | SubMb PredMode | Description | Mv Quantity | Mv Size | (Min MvSize) |
|-------------------|-------------|----------------|-------------|-------------|---------|--------------|
| OR | OR | OR | | ADD | ADD | ADD |
| 0 | 2 | 2 | B_BI_4x8 | 4 | 8 | 8 |
| 0 | 3 | 0 | BP_L0_4x4 | 4 | 8 | 4 |
| 0 | 3 | 1 | B_L1_4x4 | 4 | 8 | 4 |
| 0 | 3 | 2 | B_BI_4x4 | 8 | 8 | 8 |

3.6.4.3 Macroblock Level Rate Control

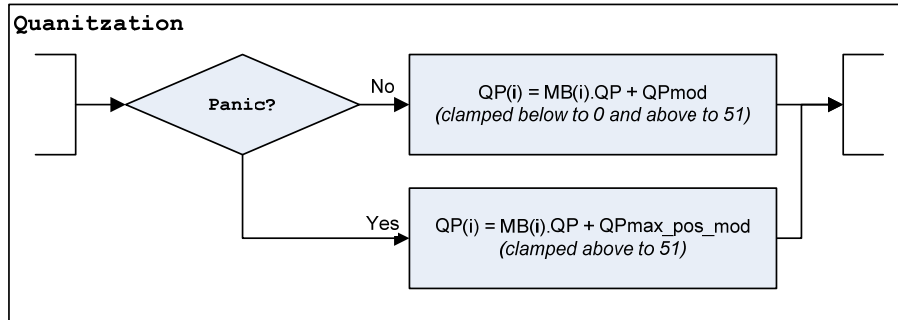
The QRC (Quantization Rate Control) unit receives data from BSP (Bit Serial Packer) and VIN (Video In) and generates adjustments to QP values across macroblocks.



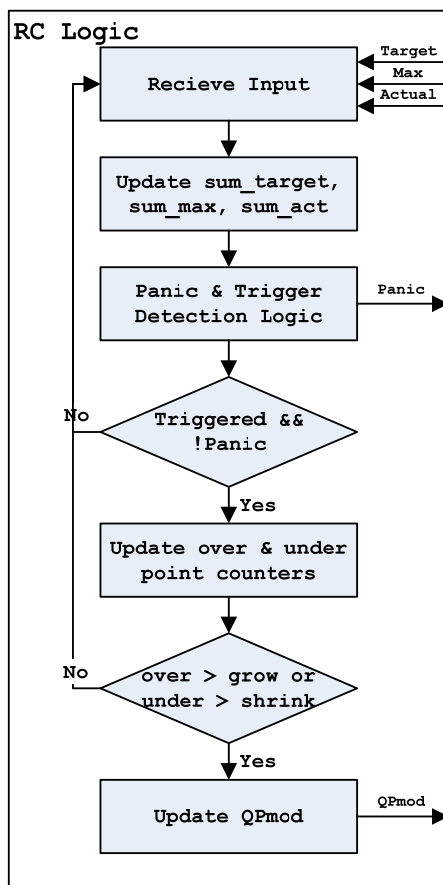
QRC can be logically partitioned into two units as shown below.



Macroblock level rate control is handled by the RC logic and the quantization logic.



The signals QPmod and panic are generated by the RC logic based on data feedback from BSP. A flowchart of the RC logic is given below.





3.6.4.3.1 Theory of Operation Overview

BSP will generate a byte estimate for each macroblock packed. Additionally, the user will specify a target and max size per macroblock. The running sum of these signals (actual, target, max) creates “curves” which are used to identify when QP adjustments are necessary (see figure below). Three more curves are symmetrically generated by QRC (upper_midpt, lower_midpt, sum_min) from target and max. The values of target and max are specified by the user will dictate the shape of these curves.

The difference between sum_actual and sum_target (called ‘bytediff’) identifies the margin of error between the target and actual sizes. The difference between the current bytediff and the previously calculated bytediff represents the rate of change in this margin over time. The sign of this rate is used to identify if the correction is trending in the appropriate direction (towards bytediff = 0).

QPmod

Each macroblock will have a requested QP (which could vary across macroblocks or remain constant). QPmod is to be added to the QP requested. QPmod will be positive when the target was under-predicted and negative when the target is over-predicted.

QPmod is incremented or decremented when internal counters (called ‘over’ and ‘under’) reach tripping points (called ‘grow’ and ‘shrink’). For each MB processed and based on which region (1-6) sum_actual falls in, various amounts of points are added to either counters. If over exceeds grow, QPmod is incremented whereas if under exceeds shrink, QPmod is decremented.

To dampen the effect of repeated changes in the same direction, an increase in resistance for that direction and decrease in resistance for the complementary direction occurs (called ‘grow_resistance’ and ‘shrink_resistance’). This resistance is added to grow or shrink, which then requires more points to trip the next correction in that direction.


The user can specify guard-bands that limit the amount QPmod can be modified. QPmod cannot exceed QPmax_pos_mod or become less than -QPmax_neg_mod_abs.

Triggering

The RC unit begins to modify QPmod occurs only when it is triggered.

Three levels of triggering exist: always, gentle, loose. Always means that RC will be active once sum_actual reaches regions 3 or 4. Gentle will trigger RC once sum_actual reaches regions 2 or 5. Loose waits to trigger RC when sum_actual reaches regions 1 or 6.

RC will deactivate (triggered = false) once sum_actual begins to track sum_target over a series of macroblocks. Specifically, the sign of the rate of change for bytediff is monitored over a window of macroblocks. When the sum of these signs over the window falls within a tolerance value (called ‘stable’), triggered will reset to false.





Panic

When enabled, panic mode will occur whenever `sum_actual` reaches region 1 and will remain so until `sum_actual` reaches region 4. When panicking, all macroblocks will be quantized with $QP = MB(n).QP + QP_{max_pos_mod}$, clamped to 51.

User Controls

This unit achieves a large flexibility by allowing the user to define various key parameters. At the per-macroblock level, the values of target and max are specified and will create various shapes of curves that `sum_actual` will be compared against.

Per-slice, the user can specify the triggering sensitivity and the tolerance required to disable the trigger. Additionally, the user can enable panic detection.

The point values assigned to each of the 6 regions are exposed to the user which allow for asymmetrical control for over and under predictions amongst other things. Additionally, the user can specify the initial values of grow and shrink along with the resistance values applied when correction is invoked.

Lastly, the maximum and minimum values for `QPmod` are also exposed to the user.

3.7 VC1 Common Commands

3.8 VC1 Decoder Commands

3.8.1 MFD_VC1_BSD_OBJECT Command

The `MFD_VC1_BSD_OBJECT` command is the only primitive command for the VC1 Decoding Pipeline. The macroblock data portion of the bitstream is loaded as indirect data object.

Before issuing a `MFD_VC1_BSD_OBJECT` command, all VC1 states of the MFD Engine need to be valid. Therefore the commands used to set these states need to have been issued prior to the issue of a `MFD_VC1_BSD_OBJECT` command.

VC1 deblock filter kernel cross the slice boundary if in the last MB row of a slice, so need to know the last MB row of a slice to disable the edge mask. There is why VC1 BSD hardware need to know the end of MB address for the current slice. As such no more phantom slice is needed for VC1, as long as the driver will program both start MB address in the current slice and the start MB address of the next slice. As a result, we can also support multiple picture state commands in between slices.

| Dword | Bits | Description |
|-------|-------|---|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | Command Opcode = MFD_VC1_BSD_OBJECT Pipeline[28:27]=2h; Opcode[26:24] = VC1 = 2h; SubOpA[23:21] = Dec = 1h; SubOpB[20:16] = 8h |



| Dword | Bits | Description |
|-------|-------|---|
| | 15:12 | Reserved : MBZ |
| | 11:0 | DWord Length (Excludes DWords 0,1) = 0002h |
| 1 | 31:22 | Reserved. MBZ |
| | 21:0 | <p>Indirect Data Length. This field provides the length in bytes of the indirect data. A value zero indicates that indirect data fetching is disabled – subsequently, the Indirect Data Start Address field is ignored.</p> <p>This field must have the same alignment as the Indirect Object Data Start Address.</p> <p>22 bits support the maximum bitstream buffer size of (e.g. for 1080i) is 3072 x120 x 68 / 8 = 3133440 bytes, which meets the requirements of the VC1 Specification.</p> <p>It includes the byte that contains the First MB Bit Offset.</p> <p>Format = U22 in bytes</p> |
| 2 | 31:29 | Reserved : MBZ |
| | 28:0 | <p>Indirect Data Start Address. This field specifies the Graphics Memory starting address of the data to be fetched into BSD Unit for processing. This pointer is relative to the MFD Indirect Object Base Address.</p> <p>It is a byte-aligned address for the VC1 bitstream data.</p> <p>Range = [0 - 512MB]</p> <p>OPEN: Supporting concealment as 0 is not allowed here [Keep this open after we deal with error handling].</p> |
| 3 | 31:24 | <p>SliceStartVertPos (Slice Vertical Position)</p> <p>This field specifies the position in y-direction of the first macroblock in the Slice in unit of macroblocks. For SecondField this value is reset to zero as opposed to the VC1 spec Ref: 9.1.2 Slice Layer.</p> |
| | 23:16 | <p>NextSliceStartVertPos (Next Slice Vertical Position)</p> <p>This field specifies the position in y-direction of the first macroblock in the next Slice in unit of macroblocks.</p> <p>This field is primarily used for error concealment. In the case that current slice is the last slice, this field should set to the height of picture (since y-direction is zero-based numbering)</p> |
| | 15:3 | Reserved : MBZ |
| | 2:0 | <p>FirstMbBitOffset (First Macroblock Bit Offset)</p> <p>This field provides the bit offset of the first macroblock of the Slice in the first byte of the input compressed bitstream.</p> <p>Format: U3</p> |

3.8.1.1 Handling Emulation Bytes

In general, VC1 BSD unit is capable of handling emulation prevention bytes. However, there is a corner case that requires host software's intervention. Host software needs to overwrite the emulation byte if it overlaps the macroblock layer decode and there is not enough information for the hardware to detect the emulation byte.

The emulation bytes might have an overlap between the picture states and the first macroblock data. If the emulation bytes are 0x00 **0x00 0x03** 0x00 and the macroblock data starts in the middle of byte1 (**0x00**),



then the host software needs to overwrite the **0x03** byte location with the previous byte (**0x00**) and change the byte offset accordingly. The hardware wouldn't know what the 1st byte was and will miss this **0x03** removal.

3.9 MPEG2 Common Commands

3.10 MPEG2 Decoder Commands

3.10.1 MFD_MPEG2_BSD_OBJECT Command (pipeline)

Different from AVC and VC1, MFD_MPEG2_BSD_OBJECT command is pipelinable. This is for performance purpose as in MPEG2 a slice is defined as a group of MBs of any size that must be within a macroblock row.

Slice header parameters are passed in as inline data and the bitstream data for the slice is passed in as indirect data. Of the inline data, slice_horizontal_position and slice_vertical_position determines the location within the destination picture of the first macroblock in the slice.

The content in this command is identical to that in the MEDIA_OBJECT command in VLD mode described in the Media Chapter.

| Dword | Bits | Description |
|-------|-------|---|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | Command Opcode = MFD_MPEG2_BSD_OBJECT Pipeline[28:27]=2h; Opcode[26:24] = MPEG2 = 3h; SubOpA[23:21] = Dec = 1h; SubOpB[20:16] = 8h |
| | 15:12 | Reserved : MBZ |
| | 11:0 | DWord Length (Excludes DWords 0,1) = 0003h |
| | 31:17 | Reserved. MBZ |
| 1 | 16:0 | Indirect Data Length. It is the length in bytes of the bitstream data for the current slice. It includes the first byte of the first macroblock and the last non-zero byte of the last macroblock in the slice. Specifically, the zero-padding bytes (if present) and the next start-code are excluded. Hardware ignores the contents after the last non-zero byte. This field is sized to support MPEG-2 MP@HL bitstream. According to Table 8-6 of ISO/IEC 13818-2, the maximum number of bits per macroblock for 4:2:0 is 4608. So the maximum slice size for MP@HL (e.g. 1080i) is $4608 * 120 / 8 = 69120$ bytes (0x10E00), which requires 17 bits. Programming Restriction: hardware has the limitation that the maximum allowed value is 0x1FFE4. As MPEG-2 spec does not post any limitation of the size of zero-padding bytes, it is possible to have a slice data with large length (including zero-padding bytes). As the data beyond 0x10E00 would only be zero bytes for a valid slice data, it is recommended that host software truncates the indirect data length to, say, 1x12000. Hardware can take care of the zero-padding bytes beyond the last non-zero byte of the slice. |
| | 31:29 | Reserved : MBZ |



| Dword | Bits | Description |
|-------|--------------|--|
| 2 | 28:0 | Indirect Data Start Address. This field specifies the Graphics Memory starting address of the data to be fetched into BSD Unit for processing. This pointer is relative to the BSD Indirect Object Base Address . Hardware ignores this field if indirect data is not present. It is a byte-aligned address for the MPEG2 VLD bitstream data |
| | 31:0 Each | Inline Data All the required Slice Header parameters and error handling settings are captured as inline data of the MPEG2_VLD_OBJECT command. It has a fixed size of 2 DWs. Its definition is described in the next section. |
| 3-4 | | |

3.10.1.1 Inline Data Description

| DWord | Bits | Description |
|-------|--|---|
| 3 | 31 | Reserved. MBZ |
| | 30:24 | Slice Horizontal Position. This 7-bit field indicates the horizontal position (in macroblock units) of the first macroblock in the slice. Format = U7 in macroblocks |
| | 23 | Reserved. MBZ |
| | 22:16 | Slice Vertical Position. This 7-bit field indicates the vertical position (in macroblock units) of the first macroblock in the slice. Format = U7 in macroblocks |
| | 15 | Reserved. MBZ |
| | 14:8 | Macroblock Count. This 7-bit field indicates the number of macroblocks in the slice, including skipped macroblocks. |
| | 7:6 | Reserved. MBZ. |
| | 5 | LastPicSlice 1 – the current Slice is the last Slice of the entire picture 0 – the current Slice is not the last Slice of current picture This bit is added to support error concealment at the end of a picture. |
| 4 | MBRowLastSlice 1 – the current Slice is the last Slice of the current MB row 0 – the current Slice is not the last Slice of MB row. This bit is added to support error concealment . | |

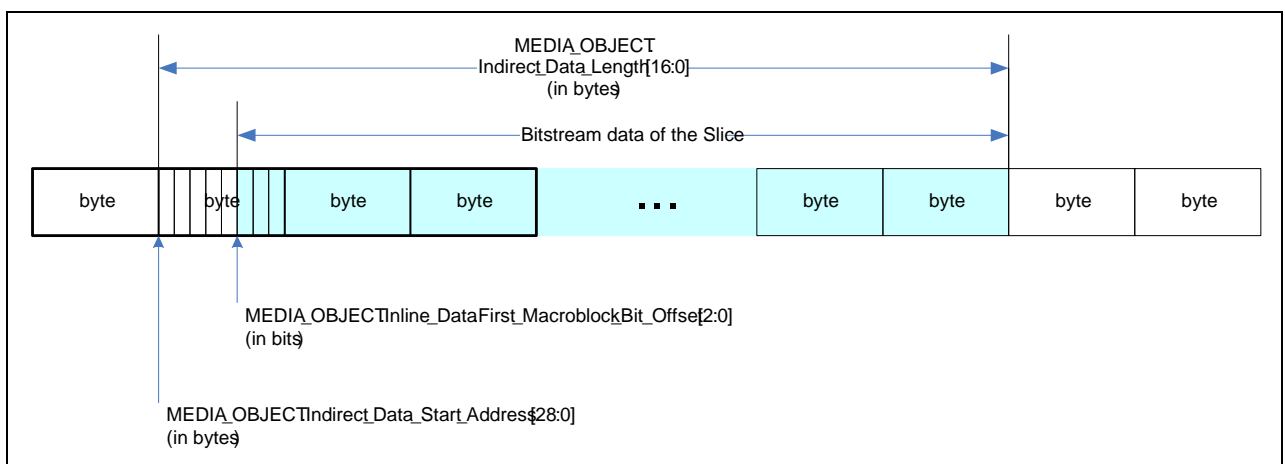
| DWord | Bits | Description |
|-------|-------|--|
| | 3 | IsLastMB 1 – the current Slice is the last Slice of the entire picture 0 – the current Slice is not the last Slice of current picture |
| | 2:0 | First Macroblock Bit Offset. This field provides the bit offset of the first macroblock in the first byte of the input bitstream. Format = U3 |
| 4 | 31:29 | Reserved. MBZ. |
| | 28:24 | Quantizer Scale Code. This field sets the quantizer scale code of the inverse quantizer. It remains in effect until changed by a decoded quantizer scale code in a macroblock. This field is decoded from the slice header by host software. Format = U5 (0 is Reserved) |
| | 23:0 | Reserved. MBZ. |

3.10.1.2 Indirect Data Description

The indirect data start address in MFD_MPEG2_BSD_OBJECT specifies the starting Graphics Memory address of the bitstream data that follows the slice header. It provides the byte address for the first macroblock of the slice. Together with the First Macroblock Bit Offset field in the inline data, it provides the bit location of the macroblock within the compressed bitstream.

The indirect data length in MFD_MPEG2_BSD_OBJECT provides the length in bytes of the bitstream data for this slice. It includes the first byte of the first macroblock and the last **non-zero** byte of the last macroblock in the slice. Specifically, the zero-padding bytes (if present) and the next start-code are excluded. Hardware ignores the contents after the last non-zero byte. Figure 3-4 illustrates these parameters for a slice data.

Figure 3-4. Indirect data buffer for a slice





3.11 MFD IT Mode Decode Commands

3.11.1 MFD_IT_OBJECT Command

All weight mode (default and implicit) are mapped to explicit mode. But the weights are come in either as explicit or implicit.

| Dword | Bits | Description |
|-------|-------|--|
| 0 | 31:29 | Command Type = PARALLEL_VIDEO_PIPE = 3h |
| | 28:16 | Command Opcode = MFD_IT_OBJECT Pipeline[28:27]=2h; Opcode[26:24] = Common = 0h; SubOpA[23:21] = Dec = 1h; SubOpB[20:16] = 9h |
| | 15:12 | Reserved : MBZ |
| | 11:0 | DWord Length (Excludes DWords 0,1) For AVC = Ch Note: Regardless of the mode, inline data must be present in this command. |
| | 31:10 | Reserved |
| 1 | 9:0 | Indirect IT-MV Data Length This field provides the length in bytes of the indirect data, which contains all the MVs for the current MB (in any partitioning and subpartitioning form). A value zero indicates that indirect data fetching is disabled – subsequently, the Indirect IT-MV Data Start Address field is ignored. This field must have the same alignment as the Indirect Object Data Start Address. AVC-IT Mode: It must be DWord aligned (since each MV is 4bytes in size) Driver has to derived this field from MVsize *4 bytes per MV. This field is only valid in AVC decoder IT mode (VC1 and MPEG uses inline MV data). Format = U10 in bytes |
| | 31:29 | Reserved : MBZ |
| 2 | 28:0 | Indirect IT-MV Data Start Address Offset This field specifies the memory starting address (offset) of the MV data to be fetched into the IT pipeline for processing. This pointer is relative to the Indirect IT-MV Object Base Address . Hardware ignores this field if indirect data is not present, i.e. the Indirect MV Data Length is set to 0. Alignment of this address depends on the mode of operation. AVC-IT Mode: It must be DWord aligned (since each MV is 4 bytes in size). This field is only valid in AVC decoder IT mode (VC1 and MPEG uses inline MV data). Range = [0 - 512MB] |
| | 31:12 | Reserved: MBZ |



| Dword | Bits | Description |
|-------|-------|---|
| 3 | 11:0 | <p>Indirect IT-COEFF Data Length</p> <p>This field provides the length in bytes of the indirect data, which contains all the non-zero coefficients for the current MB. A value zero indicates that indirect data fetching is disabled – subsequently, the Indirect IT-COEFF Data Start Address field is ignored.</p> <p>Since each IT-COEFF data is 1 DW in size, with 12 bits, this field can be extended to support up to 4:4:4 format.(256 pixel * 3 byte pixel components * 4 bytes per coeff).</p> <p>This field must be integer multiple of 16-bytes for AVC (since each coefficient is 4 bytes in size).</p> <p>This field is only valid in AVC, VC1, MPEG2 decoder IT mode.</p> <p>Format = U12 in bytes, ranging [0, 256*3*4]</p> |
| | 31:29 | Reserved: MBZ |
| 4 | 28:0 | <p>Indirect IT-COEFF Data Start Address Offset</p> <p>This field specifies the memory starting address (offset) of the coeff data to be loaded into the IT pipeline for processing. This pointer is relative to the Indirect IT-COEFF Object Base Address.</p> <p>Hardware ignores this field if indirect IT-COEFF data is not present, i.e. the Indirect IT-COEFF Data Length is set to 0.</p> <p>This field must be DW aligned, since each coefficient is 4 bytes in size.</p> <p>Driver will determine the Num of EOB 4x4/8x8 must match the block cbp flags, if not match, hardware cannot hang – add error handling.</p> <p>This field is only valid in AVC, VC1, MPEG2 decoder IT mode.</p> <p>Range = [0 - 512MB]</p> |
| | 31:6 | Reserved: MBZ |
| 5 | 5:0 | <p>Indirect IT-DBLK Control Data Length</p> <p>This field provides the length in bytes of the indirect data, which contains all the deblocker control information for the current MB (in 4x4 sub-block partitioning). A value zero indicates that indirect data fetching is disabled – subsequently, the Indirect IT-DBLK Data Start Address field is ignored.</p> <p>This field must have the same alignment as the Indirect IT-DBLK Data Start Address. It must be DWord aligned. Each Deblock Control Data record is 48 bytes or 12 DWords in size.</p> <p>This field is only valid in AVC decoder IT mode.</p> <p>Format = U6 in bytes.</p> |
| | 31:29 | Reserved: MBZ |
| 6 | 28:0 | <p>Indirect IT-DBLK Control Data Start Address Offset</p> <p>This field specifies the memory starting address (offset) of the Deblocker control data to be fetched into the IT Pipeline for processing. This pointer is relative to the Indirect IT-DBLK Object Base Address.</p> <p>Hardware ignores this field if indirect data is not present, ie. The indirect IT-DBLK Control Data Length is set to 0.</p> <p>It must be DWord aligned. Each Deblock Control Data record is 48 bytes or 12 DWords in size.</p> <p>This field is only valid in AVC decoder IT mode.</p> <p>Range = [0 - 512MB]</p> |



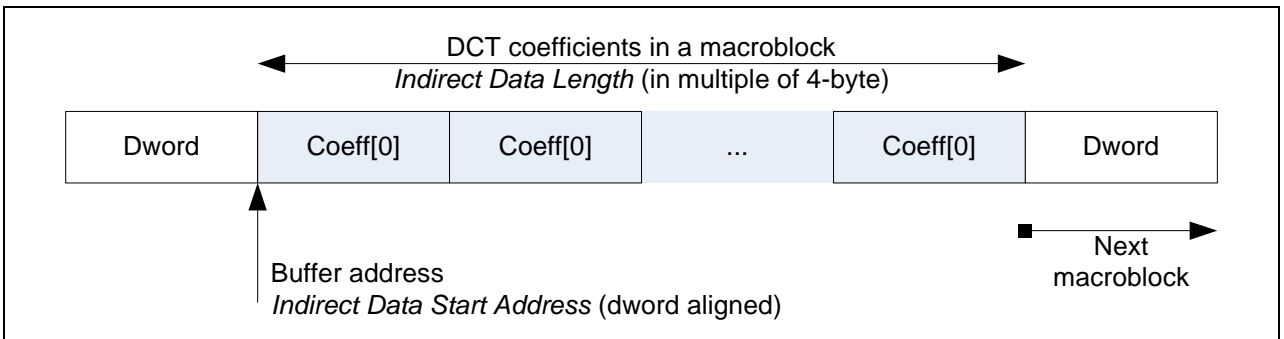
| Dword | Bits | Description |
|-------|--------------|---|
| | 31:0 Each | Inline Data (union for all 3 codecs) Includes IT, MC, IntraPred inline data as well as Deblocker control information AVC-IT Modes: Hardware interprets this data in the specified format. VC1-IT Modes: Hardware interprets this data in the specified format. MV inline MPEG2-IT Modes: Hardware interprets this data in the specified format. (IS mode) MV inline For AVC there 7 DWords of inline data, hence N is equal to 13. |
| N..7 | | |

3.11.1.1 Common Indirect IT-COEFF Data Structure

Transform-domain residual data block in AVC-IT, VC1-IT and MPEG2-IT mode follows the same data structure.

The indirect IT-COEFF data start address in MFD_IT_OBJECT command specifies the doubleword aligned address of the first non-zero DCT coefficient of the first block of the macroblock. Only the non-zero coefficients are present in the data buffer and they are packed in the 8x8 block sequence of Y0, Y1, Y2, Y3, Cb4 and Cr5, as shown in Figure 3-5. When an 8x8 block is further subdivided into 4x4 subblocks, the coefficients, if present, are organized in the subblock order. The smallest subblock division is referred to as a **transform block**. The indirect IT-COEFF data length in the command includes all the non-zero coefficients for the macroblock. It must be doubleword aligned.

Figure 3-5. Structure of the IDCT Compressed Data Buffer



Each non-zero coefficient in the indirect data buffer is contained in a doubleword-size data structure consisting of the coefficient index, end of block (EOB) flag and the fixed-point coefficient value in 2's complement form. As shown in Table 3-26, *index* is the row major 'raster' index of the coefficient **within a transform block** (please note that it is not converted to 8x8 block basis). A coefficient is a 16-bit value in 2's complement.

Table 3-26. Structure of a transform-domain residue unit

| DWord | Bit | Description |
|-------|-------|---|
| 0 | 31:16 | Transform-Domain Residual (coefficient) Value. This field contains the value of the non-zero transform-domain residual data in 2's complement. |



| DWord | Bit | Description |
|-------|------|---|
| | 15:7 | Reserved: MBZ |
| | 6:1 | <p>Index. This field specifies the raster-scan address (raw address) of the coefficient within the transform block. For a coefficient at Cartesian location (row, column) = (y, x) in a transform block of width W, Index is equal to (y * W + x). For example, coefficient at location (row, column) = (0, 0) in a 4x4 transform block has an index of 0; that at (2, 3) has an index of 2*4 + 3 = 11.</p> <p>The valid range of this field depends on the size of the transform block.</p> <p>Format = U6</p> <p>Range = [0, 63]</p> |
| | 0 | EOB (End of Block). This field indicates whether the transform-domain residue is the last one of the current transform block. |

Table 3-27. Allowed transform block dimensions per coding standard

| Transform Block Dimension | AVC | VC1 | MPEG2 |
|---------------------------|-----|-----|-------|
| 8x8 | Yes | Yes | Yes |
| 8x4 | No | Yes | No |
| 4x8 | No | Yes | No |
| 4x4 | Yes | Yes | No |

For AVC, there is intra16x16 mode, in which the DC Luma coefficients of all 4x4 sub-blocks within the current MB are sent separately in its own 4x4 Luma block. As such, only 15 coefficients remains in each of the 16 4x4 Luma blocks.

3.11.1.2 Inline Data Description in AVC-IT Mode

The Inline Data includes all the required MB decoding states, extracted primarily from the Slice Data, MB Header and their derivatives. It provides information for the following operations:

1. Inverse Quantization
2. Inverse Transform
3. Intra and inter-Prediction decoding operations
4. Internal error handling

IT Mode supports only packed MV data.

These state/parameter values may subject to change on a per-MB basis, and must be provided in each MFD_IT_OBJECT command. The values set for these variables are retained internally, until they are reset by hardware Asynchronous Reset or changed by the next MFC_AVC_PAK_OBJECT command.

The Deblocker Filter Control flags (FilterInternalEdgesFlag, FilterTopMbEdgeFlag and FilterLeftMbEdgesFlag) are generated by H/W, which are depending on MbaffFrameFlag, CurrMbAddr, PicWidthInMbs and disable_deblocking_filter_idc states.



Current MB [x,y] address is not sent, it is assumed that the H/W will keep track of the MB count and current MB position internally.

| DWord | Bit | Description |
|-------|-------|--|
| 0 | 31:24 | <p>MvQuantity</p> <p>Specify the number of MVs (in unit of motion vector, 4 bytes each) to be fetched for motion compensation operation.</p> <p>Decoder IT mode only supports packed MV format. This field specifies the exact number of MVs present for the current MB.</p> <p>For a P-Skip MB, there is still 1 MV being sent (Skip MV is sent explicitly); for a B-Direct/Skip MB, there are 2 MVs being sent.</p> <p>For an Intra-MB, MvQuantity is set to 0.</p> <p>MvQuantity = 0, signifies there is no MV indirect data for the current MB.</p> <p>This field must be set in consistent with Indirect MV Data Length, so as not to exceed its bound</p> <p>Unsigned.</p> |
| | 23:20 | Reserved MBZ |
| | 19 | <p>DcBlockCodedYFlag</p> <p>1 – the 4x4 DC-only Luma sub-block of the Intra16x16 coded MB is present; it is still possible that all DC coefficients are zero.</p> <p>0 – no 4x4 DC-only Luma sub-block is present; either not in Intra16x16 MB mode or all DC coefficients are zero.</p> |
| | 18 | <p>DcBlockCodedCbFlag</p> <p>For 4:2:0 case :</p> <p>1 – the 2x2 DC-only Chroma Cb sub-block of all coded MB (any type) is present; it is still possible that all DC coefficients are zero.</p> <p>0 – no 2x2 DC-only Chroma Cb sub-block is present; all DC coefficients are zero.</p> |
| | 17 | <p>DcBlockCodedCrFlag</p> <p>For 4:2:0 case :</p> <p>1 – the 2x2 DC-only Chroma Cr sub-block of all coded MB (any type) is present; it is still possible that all DC coefficients are zero.</p> <p>0 – no 2x2 DC-only Chroma Cr sub-block is present; all DC coefficients are zero.</p> |
| | 16 | Reserved MBZ |
| | 15 | <p>Transform8x8Flag</p> <p>0: indicates the current MB is coded with 4x4 transform and therefore the luma residuals are presented in 4x4 blocks.</p> <p>1: indicates the current MB is coded with 8x8 transform and therefore the luma residuals are presented in 8x8 blocks.</p> <p>Same as the transform_size_8x8_flag syntax element in AVC spec.</p> |



| DWord | Bit | Description |
|-------|------|---|
| | 14 | <p>MbFieldFlag</p> <p>This field specifies whether current macroblock is coded as a field or frame macroblock in MBAFF mode.</p> <p>1 = Field macroblock 0 = Frame macroblock</p> <p>This field is exactly the same as FIELD_PIC_FLAG syntax element in non-MBAFF mode. Same as the mb_field_decoding_flag syntax element in AVC spec.</p> |
| | 13 | <p>IntraMbFlag</p> <p>This field specifies whether the current macroblock is an Intra (I) macroblock.</p> <p>0 – not an intra MB 1 – is an intra MB</p> <p>I_PCM is considered as Intra MB.</p> <p>For I-picture MB (IntraPicFlag =1), this field must set to 1. This flag must be set in consistent with the interpretation of MbType (inter or intra modes).</p> |
| | 12:8 | <p>MbType</p> <p>This field carries the Macroblock Type. The meaning depends on IntraMbFlag.</p> <p>If IntraMbFlag is 1, this field is the intra macroblock type as defined in Table 3-12.</p> <p>If IntraMbFlag is 0, this field is the inter macroblock type as defined in the first two columns of Table 3-19. All macroblock types in a P Slice are mapped into the corresponding types in a B Slice. Skip and Direct modes are converted into its corresponding processing modes.</p> |
| | 7 | <p>FieldMbPolarityFlag</p> <p>This field indicates the field polarity of the current macroblock.</p> <p>Within a MbAff frame picture, this field may be different per macroblock and is set to 1 only for the second macroblock in a MbAff pair if FieldMbFlag is set. Otherwise, it is set to 0.</p> <p>Within a field picture, this field is set to 1 if the current picture is the bottom field picture. Otherwise, it is set to 0. It is a constant for the whole field picture.</p> <p>This field is only valid for MBAFF frame picture. It is reserved and set to 0 for a progressive frame picture or a field picture.</p> <p>0 = Current macroblock is a field macroblock from the top field (first in a MBAFF pair) 1 = Current macroblock is a field macroblock from the bottom field (second in a MBAFF pair)</p> |
| | 6 | <p>IsLastMB</p> <p>1 – the current MB is the last MB in the current Slice 0 – the current MB is not the last MB in the current Slice</p> |
| | 5-4 | Reserved MBZ (Intel encoder) |
| | 3:0 | Reserved MBZ |



| DWord | Bit | Description | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|-------|---|--|------|------|-------|----|-------|----|------|----|------|---|---|---|-------|----|---|---|------|---|----|----|------|---|---|---|-------|----|----|----|------|---|---|---|------|---|------|---|------|---|------|---|
| 1 | 31:16 | <p>CbpY[bit 15:0] (Coded Block Pattern Y)</p> <p>For 4x4 sub-block (when Transform8x8flag = 0 or in intra16x16) :</p> <p>16-bit cbp, one bit for each 4x4 Luma sub-block (not including the DC 4x4 Luma block in intra16x16) in a MB. The 4x4 Luma sub-blocks are numbered as</p> <table border="0"> <tr> <td>blk0</td> <td>1</td> <td>4</td> <td>5</td> <td>bit15</td> <td>14</td> <td>11</td> <td>10</td> </tr> <tr> <td>blk2</td> <td>3</td> <td>6</td> <td>7</td> <td>bit13</td> <td>12</td> <td>9</td> <td>8</td> </tr> <tr> <td>blk8</td> <td>9</td> <td>12</td> <td>13</td> <td>bit7</td> <td>6</td> <td>3</td> <td>2</td> </tr> <tr> <td>blk10</td> <td>11</td> <td>14</td> <td>15</td> <td>bit5</td> <td>4</td> <td>1</td> <td>0</td> </tr> </table> <p>The cbpY bit assignment is cbpY bit [15 – X] for sub-block_num X.</p> <p>For 8x8 block (when Transform8x8flag = 1)</p> <p>Only the lower 4 bits [3:0] are valid; the remaining upper bits [15:4] are ignored. The 8x8 Luma blocks are numbered as</p> <table border="0"> <tr> <td>blk0</td> <td>1</td> <td>bit3</td> <td>2</td> </tr> <tr> <td>blk2</td> <td>3</td> <td>bit1</td> <td>0</td> </tr> </table> <p>The cbpY bit assignment is cbpY bit [3 – X] for block_num X.</p> <p>0 in a bit – indicates the corresponding 8x8 block or 4x4 sub-block is not present (because all coefficient values are zero)</p> <p>1 in a bit – indicates the corresponding 8x8 block or 4x4 sub-block is present (although it is still possible to have all its coefficients be zero – bad coding).</p> | blk0 | 1 | 4 | 5 | bit15 | 14 | 11 | 10 | blk2 | 3 | 6 | 7 | bit13 | 12 | 9 | 8 | blk8 | 9 | 12 | 13 | bit7 | 6 | 3 | 2 | blk10 | 11 | 14 | 15 | bit5 | 4 | 1 | 0 | blk0 | 1 | bit3 | 2 | blk2 | 3 | bit1 | 0 |
| | blk0 | 1 | 4 | 5 | bit15 | 14 | 11 | 10 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | blk2 | 3 | 6 | 7 | bit13 | 12 | 9 | 8 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk8 | 9 | 12 | 13 | bit7 | 6 | 3 | 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk10 | 11 | 14 | 15 | bit5 | 4 | 1 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk0 | 1 | bit3 | 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk2 | 3 | bit1 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 15:8 | <p>VertOrigin (Vertical Origin). This field specifies the vertical origin of current macroblock in the destination picture in units of macroblocks.</p> <p>For field macroblock pair in MBAFF frame, the vertical origins for both macroblocks should be set as if they were located in corresponding field pictures. For example, for field macroblock pair originated at (16, 64) pixel location in an MBAFF frame picture, the Vertical Origin for both macroblocks should be set as 2 (macroblocks). Whether the current macroblock is the first/second (top/bottom) in a MBAFF pair is specified by FieldMbPolarityFlag.</p> <p>The macroblocks with (VertOrigin, HorzOrigin) must be delivered in the strict order as coded in the bitstream (raster order for progressive frame or field pictures and MBAFF pair order for MBAFF pictures). No gap is allowed. Otherwise, hardware behavior is undefined.</p> <p>Format = U8 in unit of macroblock.</p> | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 7:0 | <p>HorzOrigin (Horizontal Origin). This field specifies the horizontal origin of current macroblock in the destination picture in units of macroblocks.</p> <p>Format = U8 in unit of macroblock.</p> | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2 | 31:16 | <p>CbpCr (Coded Block Pattern Cr 4:2:0-only)</p> <p>Only the lower 4 bits [3:0] are valid; the remaining upper bits [15:4] are ignored (only valid for 4:2:2 and 4:4:4). The 4x4 Chroma Cr sub-blocks are numbered as</p> <table border="0"> <tr> <td>blk0</td> <td>1</td> <td>bit3</td> <td>2</td> </tr> <tr> <td>blk2</td> <td>3</td> <td>bit1</td> <td>0</td> </tr> </table> <p>The cbpCr bit assignment is cbpCr bit [3 – X] for sub-block_num X.</p> <p>0 in a bit – indicates the corresponding 4x4 sub-block is not present (because all coefficient values are zero)</p> <p>1 in a bit – indicates the corresponding 4x4 sub-block is present (although it is still possible to have all its coefficients be zero – bad coding).</p> <p>For monochrome, this field is ignored.</p> | blk0 | 1 | bit3 | 2 | blk2 | 3 | bit1 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk0 | 1 | bit3 | 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| blk2 | 3 | bit1 | 0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |



| DWord | Bit | Description | | | | | | | | |
|--------|--------------|---|------|---|------|---|------|---|------|---|
| | 15-0 | <p>CbpCb (Coded Block Pattern Cb 4:2:0-only)</p> <p>Only the lower 4 bits [3:0] are valid; the remaining upper bits [15:4] are ignored (only valid for 4:2:2 and 4:4:4). The 4x4 Chroma Cb sub-blocks are numbered as</p> <table style="margin-left: 40px;"> <tr> <td>blk0</td> <td>1</td> <td>bit3</td> <td>2</td> </tr> <tr> <td>blk2</td> <td>3</td> <td>bit1</td> <td>0</td> </tr> </table> <p>The cbpCb bit assignment is cbpCb bit [3 – X] for sub-block_num X.</p> <p>0 in a bit – indicates the corresponding 4x4 sub-block is not present (because all coefficient values are zero)</p> <p>1 in a bit – indicates the corresponding 4x4 sub-block is present (although it is still possible to have all its coefficients be zero – bad coding).</p> <p>For monochrome, this field is ignored.</p> | blk0 | 1 | bit3 | 2 | blk2 | 3 | bit1 | 0 |
| blk0 | 1 | bit3 | 2 | | | | | | | |
| blk2 | 3 | bit1 | 0 | | | | | | | |
| 3 | 31:24 | reserved MBz | | | | | | | | |
| | 23:16 | <p>QpPrimeCr</p> <p>Driver is responsible for deriving the QpPrimeCr from QpPrimeY.</p> <p>For 8-bit pixel data, QpCr is the same as QpPrimeCr, and it takes on a value in the range of 0 to 51, positive integer.</p> | | | | | | | | |
| | 15:8 | <p>QpPrimeCb</p> <p>Driver is responsible for deriving the QpPrimeCb from QpPrimeY.</p> <p>For 8-bit pixel data, QpCb is the same as QpPrimeCb, and it takes on a value in the range of 0 to 51, positive integer.</p> | | | | | | | | |
| | 7:0 | <p>QpPrimeY</p> <p>This is the per-MB QP value specified for the current MB.</p> <p>For 8-bit pixel data, QpY is the same as QpPrimeY, and it takes on a value in the range of 0 to 51, positive integer.</p> | | | | | | | | |
| 4 to 6 | 31:0 Each | <p>For intra macroblocks, definition of these fields are specified in Error! Reference source not found.</p> <p>For inter macroblocks, definition of these fields are specified in Table 3-3</p> | | | | | | | | |

3.11.1.3 Indirect Data Format in AVC-IT Mode

Indirect data in AVC-IT mode consist of Motion Vectors, Transform-domain Residue (Coefficient) and ILDB control data. All three data records have variable size. Size of each Motion Vector record is determined by the MvQuantity value as shown in Table 3-7. ILDB control record is fixed at the same size for all MBs in a picture. Coefficient data record is variable size per MB, since it may only consist of non-zero coefficients.

Each MV is represented in 4 bytes, in the form of

Lower 2 bytes : horizontal MV_x component in q-pel units

Upper 2 bytes : vertical MV_y component in q-pel units

Integer distance is measured in unit of samples in the frame or field grid position.

Chroma MVs are not sent and are derived in the H/W.



Table 3-28. Indirect MV record size in AVC-IT mode

| Macroblock Type | MVQuant |
|------------------------|----------------|
| <i>BP_L0_16x16</i> | 1 |
| <i>B_L1_16x16</i> | 1 |
| <i>B_Bi_16x16</i> | 2 |
| <i>BP_L0_L0_16x8</i> | 2 |
| <i>BP_L0_L0_8x16</i> | 2 |
| <i>B_L1_L1_16x8</i> | 2 |
| <i>B_L1_L1_8x16</i> | 2 |
| <i>B_L0_L1_16x8</i> | 2 |
| <i>B_L0_L1_8x16</i> | 2 |
| <i>B_L1_L0_16x8</i> | 2 |
| <i>B_L1_L0_8x16</i> | 2 |
| <i>B_L0_Bi_16x8</i> | 3 |
| <i>B_L0_Bi_8x16</i> | 3 |
| <i>B_L1_Bi_16x8</i> | 3 |
| <i>B_L1_Bi_8x16</i> | 3 |
| <i>B_Bi_L0_16x8</i> | 3 |
| <i>B_Bi_L0_8x16</i> | 3 |
| <i>B_Bi_L1_16x8</i> | 3 |
| <i>B_Bi_L1_8x16</i> | 3 |
| <i>B_Bi_Bi_16x8</i> | 4 |
| <i>B_Bi_Bi_8x16</i> | 4 |
| BP_8x8 | Sum |



For macroblock type of BP_8x8, MvQuant takes the sum of value MvQ[i] of the four individual 8x8 sub macroblocks.

| SubMbShape[i] | SubMbPredMode[i] | Description | MvQ[i] |
|---------------|------------------|-------------|--------|
| 0 | 0 | BP_L0_8x8 | 1 |
| 0 | 1 | B_L1_8x8 | 1 |
| 0 | 2 | B_BI_8x8 | 2 |
| 1 | 0 | BP_L0_8x4 | 2 |
| 1 | 1 | B_L1_8x4 | 2 |
| 1 | 2 | B_BI_8x4 | 4 |
| 2 | 0 | BP_L0_4x8 | 2 |
| 2 | 1 | B_L1_4x8 | 2 |
| 2 | 2 | B_BI_4x8 | 4 |
| 3 | 0 | BP_L0_4x4 | 4 |
| 3 | 1 | B_L1_4x4 | 4 |
| 3 | 2 | B_BI_4x4 | 8 |

Indirect data Deblocking Filter Control block in AVC-IT mode :

AVC Deblocker Control Data record has a fixed size for each MB in a picture and is 48 bytes or 12 Dwords in size.

| DWord | Bit | Description |
|-------|---|---|
| 0 | 31:24 | Reserved : MBZ |
| | 23 | FilterTopMbEdgeFlag |
| | 22 | FilterLeftMbEdgeFlag |
| | 21 | FilterInternal4x4EdgesFlag |
| | 20 | FilterInternal8x8EdgesFlag |
| | 19 | FieldModeAboveMbFlag |
| | 18 | FieldModeLeftMbFlag |
| | 17 | FieldModeCurrentMbFlag |
| | 16 | MbaffFrameFlag |
| | 15:8 | VertOrigin Current MB y position (address) |
| 7:0 | HorzOrigin Current MB x position (address) | |
| 1 | 31:30 | bS_h13 2-bit boundary strength for internal top horiz 4-pixel edge 3 |
| | 29:28 | bS_h12 2-bit boundary strength for internal top horiz 4-pixel edge 2 |
| | 27:26 | bS_h11 2-bit boundary strength for internal top horiz 4-pixel edge 1 |
| | 25:24 | bS_h10 2-bit boundary strength for internal top horiz 4-pixel edge 0 |



| DWord | Bit | Description |
|-------|-------|--|
| | 23:22 | bS_v33 2-bit boundary strength for internal right vert 4-pixel edge 3 |
| | 21:20 | bS_v23 2-bit boundary strength for internal right vert 4-pixel edge 2 |
| | 19:18 | bS_v13 2-bit boundary strength for internal right vert 4-pixel edge 1 |
| | 17:16 | bS_v03 2-bit boundary strength for internal right vert 4-pixel edge 0 |
| | 15:14 | bS_v32 2-bit boundary strength for internal mid vert 4-pixel edge 3 |
| | 13:12 | bS_v22 2-bit boundary strength for internal mid vert 4-pixel edge 2 |
| | 11:10 | bS_v12 2-bit boundary strength for internal mid vert 4-pixel edge 1 |
| | 9:8 | bS_v02 2-bit boundary strength for internal mid vert 4-pixel edge 0 |
| | 7:6 | bS_v31 2-bit boundary strength for internal left vert 4-pixel edge 3 |
| | 5:4 | bS_v21 2-bit boundary strength for internal left vert 4-pixel edge 2 |
| | 3:2 | bS_v11 2-bit boundary strength for internal left vert 4-pixel edge 1 |
| | 1:0 | bS_v01 2-bit boundary strength for internal left vert 4-pixel edge 0 |
| 2 | 31:28 | bS_v30_0 4-bit boundary strength for Left0 4-pixel edge 3 (MSbit is wasted) |
| | 17:24 | bS_v20_0 4-bit boundary strength for Left0 4-pixel edge 2 (MSbit is wasted) |
| | 23:20 | bS_v10_0 4-bit boundary strength for Left0 4-pixel edge 1 (MSbit is wasted) |
| | 19:16 | bS_v00_0 4-bit boundary strength for Left0 4-pixel edge 0 (MSbit is wasted) |
| | 15:14 | bS_h33 2-bit boundary strength for internal bot horiz 4-pixel edge 3 |
| | 13:12 | bS_h32 2-bit boundary strength for internal bot horiz 4-pixel edge 2 |
| | 11:10 | bS_h31 2-bit boundary strength for internal bot horiz 4-pixel edge 1 |
| | 9:8 | bS_h30 2-bit boundary strength for internal bot horiz 4-pixel edge 0 |
| | 7:6 | bS_h23 2-bit boundary strength for internal mid horiz 4-pixel edge 3 |
| | 5:4 | bS_h22 2-bit boundary strength for internal mid horiz 4-pixel edge 2 |
| | 3:2 | bS_h21 2-bit boundary strength for internal mid horiz 4-pixel edge 1 |
| | 1:0 | bS_h20 2-bit boundary strength for internal mid horiz 4-pixel edge 0 |
| 3 | 31:28 | bS_h03_0 4-bit boundary strength for Top0 4-pixel edge 3 (MSbit is wasted) |
| | 27:24 | bS_h02_0 4-bit boundary strength for Top0 4-pixel edge 2 (MSbit is wasted) |
| | 23:20 | bS_h01_0 4-bit boundary strength for Top0 4-pixel edge 1 (MSbit is wasted) |
| | 19:16 | bS_h00_0 4-bit boundary strength for Top0 4-pixel edge 0 (MSbit is wasted) |
| | 15:12 | bS_v03 4-bit boundary strength for Left1 4-pixel edge 3 (MSbit is wasted) |
| | 11:8 | bS_v02 4-bit boundary strength for Left1 4-pixel edge 2 (MSbit is wasted) |
| | 7:4 | bS_v01 4-bit boundary strength for Left1 4-pixel edge 1 (MSbit is wasted) |
| | 3:0 | bS_v00 4-bit boundary strength for Left1 4-pixel edge 0 (MSbit is wasted) |
| 4 | 31:24 | bIndexBinternal_Y Internal index B for Y |
| | 23:16 | bIndexAinternal_Y Internal index A for Y |



| DWord | Bit | Description |
|-------|-------|---|
| | 15:12 | bS_h03_1 4-bit boundary strength for Top1 4-pixel edge 3 (MSbit is wasted) |
| | 11:8 | bS_h02_1 4-bit boundary strength for Top1 4-pixel edge 2 (MSbit is wasted) |
| | 7:4 | bS_h01_1 4-bit boundary strength for Top1 4-pixel edge 1 (MSbit is wasted) |
| | 3:0 | bS_h00_1 4-bit boundary strength for Top1 4-pixel edge 0 (MSbit is wasted) |
| 5 | 31:24 | bIndexBleft1_Y |
| | 23:16 | bIndexAleft1_Y |
| | 15:8 | bIndexBleft0_Y |
| | 7:0 | bIndexAleft0_Y |
| 6 | 31:24 | bIndexBtop1_Y |
| | 23:16 | bIndexAtop1_Y |
| | 15:8 | bIndexBtop0_Y |
| | 7:0 | bIndexAtop0_Y |
| 7 | 31:24 | bIndexBleft0_Cb |
| | 23:16 | bIndexAleft0_Cb |
| | 15:8 | bIndexBinternal_Cb |
| | 7:0 | bIndexAinternal_Cb |
| 8 | 31:24 | bIndexBtop0_Cb |
| | 23:16 | bIndexAtop0_Cb |
| | 15:8 | bIndexBleft1_Cb |
| | 7:0 | bIndexAleft1_Cb |
| 9 | 31:24 | bIndexBinternal_Cr |
| | 23:16 | bIndexAinternal_Cr |
| | 15:8 | bIndexBtop1_Cb |
| | 7:0 | bIndexAtop1_Cb |
| 10 | 31:24 | bIndexBleft1_Cr |
| | 23:16 | bIndexAleft1_Cr |
| | 15:8 | bIndexBleft0_Cr |
| | 7:0 | bIndexAleft0_Cr |
| 11 | 31:24 | bIndexBtop1_Cr |
| | 23:16 | bIndexAtop1_Cr |
| | 15:8 | bIndexBtop0_Cr |
| | 7:0 | bIndexAtop0_Cr |



3.11.1.4 Inline Data Description in VC1-IT Mode

| DWord | Bit | Description | | | | | | | | | | |
|-------|---|---|-----|-------------|----|---|----|--|----|---|----|--|
| +0 | 31:28 | <p>MvFieldSelect. A bit-wise representation indicating which field in the reference frame is used as the reference field for current field. It's only used in decoding interlaced pictures. This field is valid for non-intra macroblock only.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>28</td> <td>Forward predict of current frame/field or TOP field of interlace frame, or block 0 in 4MV mode.</td> </tr> <tr> <td>29</td> <td>Backward predict of current frame/field or TOP field of interlace frame, or forward predict for block 1 in 4MV mode.</td> </tr> <tr> <td>30</td> <td>Forward predict of BOTTOM field of interlace frame, or block 2 in 4MV mode.</td> </tr> <tr> <td>31</td> <td>Backward predict of BOTTOM field of interlace frame, or forward predict for block 3 in 4MV mode.</td> </tr> </tbody> </table> <p>Each corresponding bit has the following indication. 0 = The prediction is taken from the <u>top</u> reference field. 1 = The prediction is taken from the <u>bottom</u> reference field.</p> | Bit | Description | 28 | Forward predict of current frame/field or TOP field of interlace frame, or block 0 in 4MV mode. | 29 | Backward predict of current frame/field or TOP field of interlace frame, or forward predict for block 1 in 4MV mode. | 30 | Forward predict of BOTTOM field of interlace frame, or block 2 in 4MV mode. | 31 | Backward predict of BOTTOM field of interlace frame, or forward predict for block 3 in 4MV mode. |
| | Bit | Description | | | | | | | | | | |
| | 28 | Forward predict of current frame/field or TOP field of interlace frame, or block 0 in 4MV mode. | | | | | | | | | | |
| | 29 | Backward predict of current frame/field or TOP field of interlace frame, or forward predict for block 1 in 4MV mode. | | | | | | | | | | |
| | 30 | Forward predict of BOTTOM field of interlace frame, or block 2 in 4MV mode. | | | | | | | | | | |
| 31 | Backward predict of BOTTOM field of interlace frame, or forward predict for block 3 in 4MV mode. | | | | | | | | | | | |
| 27 | Reserved. MBZ | | | | | | | | | | | |
| 26 | <p>MvFieldSelectChroma . This field specifies the polarity of reference field for chroma blocks when their motion vector is derived in Motion4MV mode for interlaced (field) picture. Non-intra macroblock only. This field is derived from MvFieldSelect. 0 = The prediction is taken from the <u>top</u> reference field. 1 = The prediction is taken from the <u>bottom</u> reference field.</p> | | | | | | | | | | | |
| 25:24 | <p>MotionType – Motion Type For frame picture, a macroblock may only be either 00 or 10. For interlace picture, a macroblock may be of any motion types. It can be 01 if and only if DctType is 1. This field is 00 if and only if IntraMacroblock is 1. 00 = Intra 01 = Field Motion. 10 = Frame Motion or no motion. Others = Reserved.</p> | | | | | | | | | | | |
| 23 | Reserved. MBZ | | | | | | | | | | | |
| | 22 | <p>MvSwitch. This field specifies whether the prediction needs to be switched from forward to backward or vice versa for single directional prediction for top and bottom fields of interlace frame B macroblocks. 0 = No directional prediction switch from top field to bottom field 1 = Switch directional prediction from top field to bottom field</p> | | | | | | | | | | |



| DWord | Bit | Description |
|-------|-------|---|
| | 21 | <p>DctType. This field specifies whether the residual data is coded as field residual or frame residual for interlaced picture. This field can be 1 only if MotionType is 00 (intra) or 01 (field motion).</p> <p>For progressive picture, this field must be set to '0', i.e. all macroblocks are frame macroblock.</p> <p>0 = Frame residual type. 1 = Field residual type.</p> |
| | 20 | <p>OverlapTransform. This field indicates whether overlap smoothing filter should be performed on I-block boundaries.</p> <p>0 = No overlap smoothing filter. 1 = Overlap smoothing filter performed.</p> |
| | 19 | <p>Motion4MV. This field indicates whether current macroblock a progressive P picture uses 4 motion vectors, one for each luminance block.</p> <p>It's only valid for progressive P-picture decoding. Otherwise, it is reserved and MBZ. For example, with MotionForward is 0, this field must also be set to 0.</p> <p>0 = 1MV-mode. 1 = 4MV-mode.</p> |
| | 18 | <p>MotionBackward. This field specifies whether the backward motion vector is active for B-picture. This field must be 0 if Motion4MV is 1 (no backward motion vector in 4MV-mode).</p> <p>0 = No backward motion vector. 1 = Use backward motion vector(s).</p> |
| | 17 | <p>MotionForward. This field specifies whether the forward motion vector is active for P and B pictures.</p> <p>0 = No forward motion vector. 1 = Use forward motion vector(s).</p> |
| | 16 | <p>IntraMacroblock. This field specifies if the current macroblock is intra-coded. When set, Coded Block Pattern is ignored and no prediction is performed (i.e., no motion vectors are used).</p> <p>For field motion, this field indicates whether the top field of the macroblock is coded as intra.</p> <p>0 = Non-intra macroblock. 1 = Intra macroblock.</p> |
| | 15:12 | <p>LumaIntra8x8Flag – Luma Intra 8x8 Flag</p> <p>This field specifies whether each of the four 8x8 luminance blocks are intra or inter coded when Motion4MV is set to 4MV-Mode.</p> <p>Each bit corresponds to one block. "0" indicates the block is inter coded and '1' indicates the block is intra coded.</p> <p>When Motion4MV is not 4MV-Mode, this field is reserved and MBZ.</p> <p>Bit 15: Y0 Bit 14: Y1 Bit 13: Y2 Bit 12: Y3</p> |



| DWord | Bit | Description |
|-------|-------|---|
| | 11:6 | <p>CBP - Coded Block Pattern</p> <p>This field specifies whether the 8x8 residue blocks in the macroblock are present or not. Each bit corresponds to one block. “0” indicates residue block isn’t present, “1” indicates residue block is present.</p> <p>Note: For each block in an intra-coded macroblock or an intra-coded block in a P macroblock in 4MV-Mode, the corresponding CBP must be 1. Subsequently, there must be at least one coefficient (this coefficient might be zero) in the indirect data buffer associated with the block (i.e. residue block must be present).</p> <p>Bit 11: Y0 Bit 10: Y1 Bit 9: Y2 Bit 8: Y3 Bit 7: Cb4 Bit 6: Cr5</p> |
| | 5 | <p>ChromaIntraFlag - Derived Chroma Intra Flag</p> <p>This field specifies whether the chroma blocks should be treated as intra blocks based on motion vector derivation process in 4MV mode.</p> <p>0 = Chroma blocks are not coded as intra. 1 = Chroma blocks are coded as intra</p> |
| | 4 | <p>LastRowFlag – Last Row Flag</p> <p>This field indicates that the current macroblock belongs to the last row of the picture. This field may be used by the kernel to manage pixel output when overlap transform is on.</p> <p>0 = Not in the last row 1 = In the last row</p> |
| | 3 | LastMBInRow – This field indicates the last MB in row flag. |
| | 2:0 | Reserved. MBZ |
| +1 | 32:26 | Reserved. MBZ |
| | 15:8 | <p>VertOrigin - Vertical Origin</p> <p>In unit of macroblocks relative to the current picture (frame or field).</p> |
| | 7:0 | <p>HorzOrigin - Horizontal Origin</p> <p>In unit of macroblocks.</p> |
| +2 | 31:16 | MotionVector[0].Vert |
| | 15:0 | MotionVector[0].Horz |
| +3 | 31:0 | MotionVector[1] |
| +4 | 31:0 | MotionVector[2] |
| +5 | 31:0 | MotionVector[3] |
| +6 | 31:0 | <p>MotionVectorChroma</p> <p>This field is not valid for a field motion in an interlaced frame picture where 4 MVs for chroma blocks.</p> <p>Notes: This field is derived from MotionVector[3:0] as described in the following section.</p> |



| DWord | Bit | Description | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|-------|-----------------------------|--|---------------------------------------|--|--|------------|------------|--|------|---------|-------|-------|-------|-------|----|------------------------|-----|------------|------------|------------|----|---------------------------|-----|-----|------------|------------|----|---------------------------|-----|-----|------------|------------|----|----------------------------|-----|-----|-----|-----|
| +7 | 32:24 | Subblock Code for Y3 The following subblock coding definition applies to all 6 subblock coding bytes. Bits 7:6 are reserved. <table border="1" style="margin-left: 20px;"> <thead> <tr> <th colspan="2">Subblock Partitioning (Bits [1:0])</th> <th colspan="4">Subblock Present (0 means not present, 1 means present)</th> </tr> <tr> <th>Code</th> <th>Meaning</th> <th>Bit 2</th> <th>Bit 3</th> <th>Bit 4</th> <th>Bit 5</th> </tr> </thead> <tbody> <tr> <td>00</td> <td>Single 8x8 block (sb0)</td> <td>Sb0</td> <td>Don't care</td> <td>Don't care</td> <td>Don't care</td> </tr> <tr> <td>01</td> <td>Two 8x4 subblocks (sb0-1)</td> <td>Sb0</td> <td>Sb1</td> <td>Don't care</td> <td>Don't care</td> </tr> <tr> <td>10</td> <td>Two 4x8 subblocks (sb0-1)</td> <td>Sb0</td> <td>Sb1</td> <td>Don't care</td> <td>Don't care</td> </tr> <tr> <td>11</td> <td>Four 4x4 subblocks (sb0-3)</td> <td>Sb0</td> <td>Sb1</td> <td>Sb2</td> <td>Sb3</td> </tr> </tbody> </table> | Subblock Partitioning (Bits [1:0]) | | Subblock Present (0 means not present, 1 means present) | | | | Code | Meaning | Bit 2 | Bit 3 | Bit 4 | Bit 5 | 00 | Single 8x8 block (sb0) | Sb0 | Don't care | Don't care | Don't care | 01 | Two 8x4 subblocks (sb0-1) | Sb0 | Sb1 | Don't care | Don't care | 10 | Two 4x8 subblocks (sb0-1) | Sb0 | Sb1 | Don't care | Don't care | 11 | Four 4x4 subblocks (sb0-3) | Sb0 | Sb1 | Sb2 | Sb3 |
| | | Subblock Partitioning (Bits [1:0]) | | Subblock Present (0 means not present, 1 means present) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | Code | Meaning | Bit 2 | Bit 3 | Bit 4 | Bit 5 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 00 | Single 8x8 block (sb0) | Sb0 | Don't care | Don't care | Don't care | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 01 | Two 8x4 subblocks (sb0-1) | Sb0 | Sb1 | Don't care | Don't care | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 10 | Two 4x8 subblocks (sb0-1) | Sb0 | Sb1 | Don't care | Don't care | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 11 | Four 4x4 subblocks (sb0-3) | Sb0 | Sb1 | Sb2 | Sb3 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 23:16 | Subblock Code for Y2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 15:8 | Subblock Code for Y1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 7:0 | Subblock Code for Y0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| +8 | 31:16 | Reserved. MBZ | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 15:8 | Subblock Code for Cr | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 7:0 | Subblock Code for Cb | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| +9 | 31:24 | ILDB control data for block Y3 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 23:16 | ILDB control data for block Y2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 15:8 | ILDB control data for block Y1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| S | 7:0 | ILDB control data for block Y0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 31:16 | Reserved | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 15:8 | ILDB control data for Cr block | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | 7:0 | ILDB control data for Cb block | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

3.11.1.5 Indirect Data Format in VC1-IT Mode

VC1-IT mode only contains IT-COEFF indirect data which is described in Section 3.11.1.1.

3.11.1.6 Inline Data Description in MPEG2-IT Mode

The content in this command is similar to that in the MEDIA_OBJECT command in IS mode described in the Media Chapter.

Each MFD_IT_OBJECT command corresponds to the processing of one macroblock. Macroblock parameters are passed in as inline data and the non-zero DCT coefficient data for the macroblock is passed in as indirect data.

Error! Reference source not found. depicts the inline data format. Inline data starts at dword 7 of MFD_IT_OBJECT command. There are 7 dwords total.



Table 3-29. Inline data in MPEG2-IT Mode

| DWord | Bit | Description | | | | | | | | | | | | | | | | | | | | |
|-------|--|--|--|---|--|---------------------------------|----------|----------|------|-------|-------|------|-------|------|------|------------|------------|---|----|---|---|---|
| +0 | 31:28 | <p>Motion Vertical Field Select. A bit-wise representation of a long [2][2] array as defined in §6.3.17.2 of the <i>ISO/IEC 13818-2</i> (see also §7.6.4).</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>MVector [r]</th> <th>MVector [s]</th> <th>MotionVerticalFieldSelect Index</th> </tr> </thead> <tbody> <tr> <td>28</td> <td>0</td> <td>0</td> <td>0</td> </tr> <tr> <td>29</td> <td>0</td> <td>1</td> <td>1</td> </tr> <tr> <td>30</td> <td>1</td> <td>0</td> <td>2</td> </tr> <tr> <td>31</td> <td>1</td> <td>1</td> <td>3</td> </tr> </tbody> </table> <p>Format = MC_MotionVerticalFieldSelect. 0 = The prediction is taken from the <u>top</u> reference field. 1 = The prediction is taken from the <u>bottom</u> reference field.</p> | Bit | MVector [r] | MVector [s] | MotionVerticalFieldSelect Index | 28 | 0 | 0 | 0 | 29 | 0 | 1 | 1 | 30 | 1 | 0 | 2 | 31 | 1 | 1 | 3 |
| | Bit | MVector [r] | MVector [s] | MotionVerticalFieldSelect Index | | | | | | | | | | | | | | | | | | |
| | 28 | 0 | 0 | 0 | | | | | | | | | | | | | | | | | | |
| | 29 | 0 | 1 | 1 | | | | | | | | | | | | | | | | | | |
| 30 | 1 | 0 | 2 | | | | | | | | | | | | | | | | | | | |
| 31 | 1 | 1 | 3 | | | | | | | | | | | | | | | | | | | |
| 27 | Reserved (was Second Field) | | | | | | | | | | | | | | | | | | | | | |
| 26 | Reserved. (HWMC mode) | | | | | | | | | | | | | | | | | | | | | |
| 25:24 | 25:24 | <p>Motion Type. When combined with the destination picture type (field or frame) this Motion Type field indicates the type of motion to be applied to the macroblock. See <i>ISO/IEC 13818-2</i> §6.3.17.1, Tables 6-17, 6-18. In particular, the device supports dual-prime motion prediction (11) in both frame and field picture type.</p> <p>Format = MC_MotionType</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Destination = Frame Picture_Structure = 11</th> <th>Destination = Field Picture_Structure != 11</th> </tr> </thead> <tbody> <tr> <td>'00'</td> <td>Reserved</td> <td>Reserved</td> </tr> <tr> <td>'01'</td> <td>Field</td> <td>Field</td> </tr> <tr> <td>'10'</td> <td>Frame</td> <td>16x8</td> </tr> <tr> <td>'11'</td> <td>Dual-Prime</td> <td>Dual-Prime</td> </tr> </tbody> </table> | Value | Destination = Frame Picture_Structure = 11 | Destination = Field Picture_Structure != 11 | '00' | Reserved | Reserved | '01' | Field | Field | '10' | Frame | 16x8 | '11' | Dual-Prime | Dual-Prime | | | | | |
| | Value | Destination = Frame Picture_Structure = 11 | Destination = Field Picture_Structure != 11 | | | | | | | | | | | | | | | | | | | |
| | '00' | Reserved | Reserved | | | | | | | | | | | | | | | | | | | |
| | '01' | Field | Field | | | | | | | | | | | | | | | | | | | |
| '10' | Frame | 16x8 | | | | | | | | | | | | | | | | | | | | |
| '11' | Dual-Prime | Dual-Prime | | | | | | | | | | | | | | | | | | | | |
| 23:22 | Reserved. (Scan method) | | | | | | | | | | | | | | | | | | | | | |
| 21 | <p>DCT Type. This field specifies the DCT type of the current macroblock. The kernel should ignore this field when processing Cb/Cr data. See <i>ISO/IEC 13818-2</i> §6.3.17.1. This field is zero if Coded Block Pattern is also zero (no coded blocks present).</p> <p>0 = MC_FRAME_DCT (Macroblock is frame DCT coded). 1 = MC_FIELD_DCT (Macroblock is field DCT coded).</p> | | | | | | | | | | | | | | | | | | | | | |
| 20 | Reserved (was Overlap Transform - H261 Loop Filter). | | | | | | | | | | | | | | | | | | | | | |
| 18 | <p>Macroblock Motion Backward. This field specifies if the backward motion vector is active. See <i>ISO/IEC 13818-2</i> Tables B-2 through B-4.</p> <p>0 = No backward motion vector. 1 = Use backward motion vector(s).</p> | | | | | | | | | | | | | | | | | | | | | |
| 17 | <p>Macroblock Motion Forward. This field specifies if the forward motion vector is active. See <i>ISO/IEC 13818-2</i> Tables B-2 through B-4.</p> <p>0 = No forward motion vector. 1 = Use forward motion vector(s).</p> | | | | | | | | | | | | | | | | | | | | | |



| DWord | Bit | Description |
|-------|-------|---|
| | 16 | Macroblock Intra Type. This field specifies if the current macroblock is intra-coded. When set, Coded Block Pattern is ignored and no prediction is performed (i.e., no motion vectors are used). See <i>ISO/IEC 13818-2</i> Tables B-2 through B-4. 0 = Non-intra macroblock. 1 = Intra macroblock. |
| | 15:12 | Reserved : MBZ |
| | 11:6 | Coded Block Pattern. This field specifies whether blocks are present or not. Format = 6-bit mask. Bit 11: Y0 Bit 10: Y1 Bit 9: Y2 Bit 8: Y3 Bit 7: Cb4 Bit 6: Cr5 |
| | 5:4 | Reserved. (Quantization Scale Code) |
| | 3 | LastMBInRow – This field indicates the last MB in each row. |
| | 2:0 | Reserved: MBZ |
| +1 | 31:16 | Reserved : MBZ |
| | 15:8 | VertOrigin - Vertical Origin In unit of macroblocks relative to the current picture (frame or field). |
| | 7:0 | HorzOrigin - Horizontal Origin In unit of macroblocks. |
| +2 | 31:16 | Motion Vectors – Field 0, Forward, Vertical Component. Each vector component is a 16-bit two's-complement value. The vector is relative to the current macroblock location. According to <i>ISO/IEC 13818-2</i> Table 7-8, the valid range of each vector component is [-2048, +2047.5], implying a format of s11.1. However, it should be noted that motion vector values are sign extended to 16 bits. |
| | 15:0 | Motion Vectors – Field 0, Forward, Horizontal Component |
| +3 | 31:16 | Motion Vectors – Field 0, Backward, Vertical Component |
| | 15:0 | Motion Vectors – Field 0, Backward, Horizontal Component |
| +4 | 31:16 | Motion Vectors – Field 1, Forward, Vertical Component |
| | 15:0 | Motion Vectors – Field 1, Forward, Horizontal Component |
| +5 | 31:16 | Motion Vectors – Field 1, Backward, Vertical Component |
| | 15:0 | Motion Vectors – Field 1, Backward, Horizontal Component |

3.11.1.7 Indirect Data Format in MPEG2-IT Mode

MPEG2-IT mode only contains IT-COEFF indirect data which is described in Section 3.11.1.1.



3.12 AVC Encoder MBAFF Support

1. Algorithm

Prediction of current macroblock motion vector is possible from neighboring macroblocks $mbAddrA/mbAddrD/mbAddrB/mbAddrC/mbAddrA+1/mbAddrD+1/mbAddrB+1/mbAddrC+1$. The selection of these macroblocks depends on coding type (field/frame) of current macroblock pair and the coding of neighbouring macroblock pair. Following is a generic diagram depicting naming conventions used for neighbouring macroblocks. Selection of these mb pairs described in detail in following sections.

1.1. Selection of Top Left MB pair:

The selection of Top Left MB pair depends on coding type of current and also top left macroblock pair. Following diagram shows the mapping to be used in MPC unit for the selection of the Top Left MB (D or D+1 macroblock).

1.2. Selection of Left MB pair

The selection of Left MB pair depends on coding type of current and also left macroblock pair. Following diagram shows the mapping to be used in MPC unit for the selection of the Left MB (A or A+1 macroblock).

1.3. Selection of Top MB pair

The selection of Top MB pair depends on coding type of current and also top macroblock pair. Following diagram shows the mapping to be used in MPC unit for the selection of the Top MB (B or B+1 macroblock).

1.4. Selection of Top Right MB pair

The selection of Top Right MB pair depends on coding type of current and also top right macroblock pair. Following diagram shows the mapping to be used in MPC unit for the selection of the Top Right MB (C or C+1 macroblock).

1.5. Motion Vector and refIdx Scaling

Motion vectors and refence index of neighbouring macroblocks ($mbAddrA/mbAddrB/mbAddrC/mbAddrD$) should be scaled before using them into prediction equations. Again the scaling depends on coding type of current and neighbouring macroblock pair which is described as follows,

- If the current macroblock is a field macroblock and the macroblock $mbAddrN$ is a frame macroblock

$$mvLXN[1] = mvLXN[1] / 2 \quad (8-214)$$

$$refIdxLXN = refIdxLXN * 2 \quad (8-215)$$

- Otherwise, if the current macroblock is a frame macroblock and the macroblock $mbAddrN$ is a field macroblock



$$mvLXN[1] = mvLXN[1] * 2 \quad (8-216)$$

$$refIdxLXN = refIdxLXN / 2 \quad (8-217)$$

- Otherwise, the vertical motion vector component $mvLXN[1]$ and the reference index $refIdxLXN$ remain unchanged.

3.13 Decoder StreamOut Mode Data Structure Definition

When StreamOut is enabled, per MB intermediated decoded data (MVs, mb_type, MB qp, etc.) are sent to the memory in a fixed record format (and of fixed size). The per-MB records must be written in a strict raster order and with no gap (i.e. every MB regardless of its mb_type and slice type, must have an entry in the StreamOut buffer). Therefore, the consumer of the StreamOut data can offset into the StreamOut Buffer (**StreamOut Data Destination Base Address**) using individual MB addresses.

A StreamOut Data record format is detailed as follows :

The inline data content of Dwords 4 to 6 is defined either for intra prediction or for inter prediction, but not both.

| DWord | Bit | Description |
|-------|-------|--|
| | 23 | Reserved MBZ |
| | 22-20 | MvFormat (see PAK_OBJECT) |
| | 19:17 | CodedPatternDC (for AVC only, 111b for others) The field indicates whether DC coefficients are sent.. 1 bit each for Y, U and V. |
| | 16 | Reserved MBZ |
| | 15 | Transform8x8Flag When it is set to 0, the current MB uses 4x4 transform. When it is set to 1, the current MB uses 8x8 transform. The transform_size_8x8_flag syntax element, if present in the output bitstream, is the same as this field. However, whether transform_size_8x8_flag is present or not in the output bitstream depends on several conditions: This field is only allowed to be set to 1 for two conditions: <ul style="list-style-type: none"> • It must be 1 if IntraMbFlag = INTRA and IntraMbMode = INTRA_8x8 • It may be 1 if IntraMbFlag = INTER and there is no sub partition size less than 8x8 Otherwise, this field must be set to 0. 0: 4x4 integer transform 1: 8x8 integer transform |



| DWord | Bit | Description |
|-------|-------|---|
| | 14 | <p>MbFieldFlag</p> <p>This field specifies whether current macroblock is coded as a field or frame macroblock in MBAFF mode.</p> <p>This field is exactly the same as FIELD_PIC_FLAG syntax element in non-MBAFF mode. Same as the mb_field_decoding_flag syntax element in AVC spec.</p> <p>0 = Frame macroblock 1 = Field macroblock</p> |
| | 13 | <p>IntraMbFlag</p> <p>This field specifies whether the current macroblock is an Intra (I) macroblock. I_PCM is considered as Intra MB.</p> <p>For I-picture MB (IntraPicFlag =1), this field must be set to 1.</p> <p>This flag must be set in consistent with the interpretation of MbType (inter or intra modes).</p> <p>0: INTER (inter macroblock)</p> <p>1: INTRA (intra macroblock)</p> |
| | 12:8 | <p>MbType5Bits</p> <p>This field is encoded to match with the best macroblock mode determined as described in the next section. It follows AVC encoding for inter and intra macroblocks.</p> |
| | 7:6 | <p>Reserved MBZ</p> |
| | 5:4 | <p>IntraMbMode</p> <p>This field is provided to carry information partially overlapped with MbType.</p> <p>This field is only valid if IntraMbFlag = INTRA, otherwise, it is ignored by hardware..</p> |
| | 2:0 | <p>InterMbMode</p> <p>This field is provided to carry redundant information as that in MbType. It also carries additional information such as skip.</p> <p>This field is only valid if IntraMbFlag =INTER, otherwise, it is ignored by hardware.</p> |
| 1 | 31:16 | <p>MbYCnt (Vertical Origin). This field specifies the vertical origin of current macroblock in the destination picture in units of macroblocks.</p> <p>Format = U8 in unit of macroblock.</p> |
| | 15:0 | <p>MbXCnt (Horizontal Origin). This field specifies the horizontal origin of current macroblock in the destination picture in units of macroblocks.</p> <p>Format = U8 in unit of macroblock.</p> |
| 2 | 31:24 | <p>NzCoefCountY3</p> <p>The number of non-zero coefficients involved in the 4th 8x8 sub-block area for Luma. Maximal number of AC coefficients allowed in a 4x4 transform.sub-block (For 8x8 transforms the number will be scaled up by 4x accordingly).</p> <p>It takes on a value in the range of 0 to 15 for a 4x4 sub-block, and in the range of 0 to 63 for an 8x8 block.</p> |
| | 23:16 | <p>NzCoefCountY2</p> <p>The number of non-zero coefficients involved in the 3rd 8x8 sub-block area for Luma. Maximal number of AC coefficients allowed in a 4x4 transform.sub-block (For 8x8 transforms the number will be scaled up by 4x accordingly).</p> <p>It takes on a value in the range of 0 to 15 for a 4x4 sub-block, and in the range of 0 to 63 for an 8x8 block.</p> |



| DWord | Bit | Description |
|-------|----------------|---|
| | 15:8 | <p>NzCoefCountY1</p> <p>The number of non-zero coefficients involved in the 2nd 8x8 sub-block area for Luma. Maximal number of AC coefficients allowed in a 4x4 transform.sub-block (For 8x8 transforms the number will be scaled up by 4x accordingly). It takes on a value in the range of 0 to 15 for a 4x4 sub-block, and in the range of 0 to 63 for an 8x8 block.</p> |
| | 7:0 | <p>NzCoefCountY0</p> <p>The number of non-zero coefficients involved in the 1st 8x8 sub-block area for Luma. Maximal number of AC coefficients allowed in a 4x4 transform.sub-block (For 8x8 transforms the number will be scaled up by 4x accordingly). It takes on a value in the range of 0 to 15 for a 4x4 sub-block, and in the range of 0 to 63 for an 8x8 block.</p> |
| 3 | 31:28 AVC | <p>Skip8x8Pattern (AVC)</p> <p>This field indicates whether each of the four 8x8 sub macroblocks is using the predicted MVs and will not be explicitly coded in the bitstream (the sub macroblock will be coded as direct mode). It contains four 1-bit subfields, corresponding to the 4 sub macroblocks in sequential order. The whole macroblock may be actually coded as B_Direct_16x16 or B_Skip, according to the macroblock type conversion rules described in a later sub section. This field is only valid for a B slice. It is ignored by hardware for a P slice. Hardware also ignores this field for an intra macroblock.</p> <p>0 in a bit – Corresponding MVs are sent in the bitstream 1 in a bit – Corresponding MVs are not sent in the bitstream</p> |
| | 31 MPEG2 | <p>EndSliceFlag (MPEG2)</p> <p>1 – the current MB is the last MB in the current Slice 0 – the current MB is not the last MB in the current SliceReserved MBZ.</p> |
| | 30 MPEG2 | <p>NewSliceFlag (MPEG2)</p> <p>1 – the current MB is the first MB in the current Slice 0 – the current MB is not the first MB in the current SliceReserved MBZ.</p> |
| | 29:28 MPEG2 | <p>MbScanMethod (MPEG2)</p> <p>Same value as in picture/slice parameter</p> |
| | 27 | <p>EnableCoeffClamp</p> <p>1 = the magnitude of coefficients of the current MB will be clamped based on the clamping matrix after quantization 0 = no clamping</p> |
| | 26 | <p>LastMbFlag</p> <p>1 – the current MB is the last MB in the current call (slice for AVC, slice group for MPEG2) 0 – the current MB is not the last MB in the current call.</p> |
| | 25 | <p>SkipMbConvDisable</p> <p>This is a per-MB level control to enable and disable skip conversion. This field is ORed with SkipConvDisable field. This field is only valid for a P or B slice. It must be zero for other slice types. Rules are provided in Section 3.6.4.1.6</p> <p>0 - Enable skip type conversion for the current macroblock 1 - Disable skip type conversion for the current macroblock</p> |
| | 24 | <p>Reserved MBZ</p> |



| DWord | Bit | Description |
|---------------|--------------|--|
| | 23:16 | NzCoefCountV Maximal number of AC coefficients allowed in a 4x4 transform.sub-block (For 8x8 transforms the number will be scaled up by 4x accordingly). It takes on a value in the range of 0 to 15 for a 4x4 sub-block, and in the range of 0 to 63 |
| | 15:8 | NzCoefCountU Maximal number of AC coefficients allowed in a 4x4 transform.sub-block (For 8x8 transforms the number will be scaled up by 4x accordingly). It takes on a value in the range of 0 to 15 for a 4x4 sub-block, and in the range of 0 to 63 |
| | 7 | mbz (AVC) / QScaleType (MPEG2) |
| | 6:0 | QpPrimeY (AVC) / QScaleCode (MPEG2) The luma quantization index. This is the per-MB QP value specified for the current MB. For SNB, it will always return 0xFF (unknown). |
| 4 to 6 | 31:0 Each | For intra macroblocks, definition of these fields are specified in Table 1 1 For inter macroblocks, definition of these fields are specified in Table 1 2 |
| 7 | 31:0 | Reserved MBz. |
| Inter cases | | |
| 8 | 31:16 | MV[0][0].y – y-component of the forward motion vector of the 1 st 8x8 or 1 st 4x4 subblock |
| | 15:0 | MV[0][0].x – x-component of the forward motion vector of the 1 st 8x8 or 1 st 4x4 subblock |
| 9 | 31:0 | MV[1][0] – the backward motion vector of the 1 st 8x8 or 1 st 4x4 subblock |
| 10 | 31:0 | MV[0][1] – the forward motion vector of the 2 nd 8x8 or 4 th 4x4 subblock |
| 11 | 31:0 | MV[1][1] – the backward motion vector of the 2 nd 8x8 or 4 th 4x4 subblock |
| 12 | 31:0 | MV[0][2] – the forward motion vector of the 3 rd 8x8 or 8 th 4x4 subblock |
| 13 | 31:0 | MV[1][2] – the backward motion vector of the 3 rd 8x8 or 8 th 4x4 subblock |
| 14 | 31:0 | MV[0][3] – the forward motion vector of the 4 th 8x8 or 12 th 4x4 subblock |
| 15 | 31:0 | MV[1][3] – the backward motion vector of the 4 th 8x8 or 12 th 4x4 subblock |
| Intra Cases : | | |
| 8 to 15 | 31:0 | Reserved MBZ |

Inline data subfields for an Intra Macroblock

| DWord | Bit | Description |
|-------|-------|---|
| 4 | 31:16 | LumaIntraPredModes[1] Specifies the Luma Intra Prediction mode for four 4x4 sub-block of a MB, 4-bit each. AVC : See the bit assignment table later in this section. VC1 : MBZ. MPEG2 : MBZ. |



| DWord | Bit | Description |
|-------------------|-------|---|
| 5 AVC INTRA | 15:0 | <p>LumaIndraPredModes[0] Specifies the Luma Intra Prediction mode for four 4x4 sub-block, four 8x8 block or one intra16x16 of a MB. 4-bit per 4x4 sub-block (Transform8x8Flag=0, Mbtype=0 and intraMbFlag=1) or 8x8 block (Transform8x8Flag=1, Mbtype=0, MbFlag=1), since there are 9 intra modes. 4-bit for intra16x16 MB (Transform8x8Flag=0, Mbtype=1 to 24 and intraMbFlag=1), but only the LSBit[1:0] is valid, since there are only 4 intra modes. AVC : See the bit assignment table later in this section. VC1 : MBZ. MPEG2 : MBZ.</p> |
| | 31:16 | <p>LumaIndraPredModes[3] Specifies the Luma Intra Prediction mode for four 4x4 sub-block of a MB, 4-bit each. AVC : See the bit assignment table later in this section. VC1 : MBZ. MPEG2 : MBZ.</p> |
| | 15:0 | <p>LumaIndraPredModes[2] Specifies the Luma Intra Prediction mode for four 4x4 sub-block of a MB, 4-bit each. AVC : See the bit assignment later in this section. VC1 : MBZ. MPEG2 : MBZ.</p> |
| 6 | 31:8 | <p>Reserved (Reserved for encoder turbo mode IntraResidueDataSize, when this is not 0, optional residue data are provided to the PAK; Reserved for decoder)</p> |



| DWord | Bit | Description | | | | | | | | | | | | | | | | |
|-------|--|--|------|-------------------------------------|---|---|---|---|---|--|---|---|---|---|---|---|-----|--|
| | 7:0 | <p>MbIntraStruct</p> <p>The IntraPredAvailFlags[4:0] have already included the effect of the constrained_intra_pred_flag. See the diagram later for the definition of neighbors position around the current MB or MB pair (in MBAFF mode).</p> <p>1 – IntraPredAvailFlagX, indicates the values of samples of neighbor X can be used in intra prediction for the current MB.</p> <p>0 – IntraPredAvailFlagX, indicates the values of samples of neighbor X is not available for intra prediction of the current MB.</p> <p>IntraPredAvailFlag-A and -E can only be different from each other when constrained_intra_pred_flag is equal to 1 and mb_field_decoding_flag is equal to 1 and the value of the mb_field_decoding_flag for the macroblock pair to the left of the current macroblock is equal to 0 (which can only occur when MbaffFrameFlag is equal to 1).</p> <p>IntraPredAvailFlag-F is used only if</p> <ul style="list-style-type: none"> • it is in MBAFF mode, i.e. MbaffFrameFlag = 1, • the current macroblock is of frame type, i.e. MbFieldFag = 0, and • the current macroblock type is Intra8x8, i.e. IntraMbFlag = INTRA, IntraMbMode = INTRA_8x8, and Transform8x8Flag = 1. <p>In any other cases IntraPredAvailFlag-A shall be used instead.</p> <table border="1" data-bbox="550 957 1430 1344"> <thead> <tr> <th>Bits</th> <th>IntraPredAvailFlags[4:0] Definition</th> </tr> </thead> <tbody> <tr> <td>7</td> <td>IntraPredAvailFlagF – F (Left 8th row (-1,7) neighbor)</td> </tr> <tr> <td>6</td> <td>IntraPredAvailFlagA – A (Left neighbor top half)</td> </tr> <tr> <td>5</td> <td>IntraPredAvailFlagE – E (Left neighbor bottom half)</td> </tr> <tr> <td>4</td> <td>IntraPredAvailFlagB – B (Top neighbor)</td> </tr> <tr> <td>3</td> <td>IntraPredAvailFlagC – C (Top right neighbor)</td> </tr> <tr> <td>2</td> <td>IntraPredAvailFlagD – D (Top left corner neighbor)</td> </tr> <tr> <td>1:0</td> <td>ChromalIntraPredMode – 2 bits to specify 1 of 4 chroma intra prediction mode, see the table in later section.</td> </tr> </tbody> </table> | Bits | IntraPredAvailFlags[4:0] Definition | 7 | IntraPredAvailFlagF – F (Left 8 th row (-1,7) neighbor) | 6 | IntraPredAvailFlagA – A (Left neighbor top half) | 5 | IntraPredAvailFlagE – E (Left neighbor bottom half) | 4 | IntraPredAvailFlagB – B (Top neighbor) | 3 | IntraPredAvailFlagC – C (Top right neighbor) | 2 | IntraPredAvailFlagD – D (Top left corner neighbor) | 1:0 | ChromalIntraPredMode – 2 bits to specify 1 of 4 chroma intra prediction mode, see the table in later section. |
| Bits | IntraPredAvailFlags[4:0] Definition | | | | | | | | | | | | | | | | | |
| 7 | IntraPredAvailFlagF – F (Left 8 th row (-1,7) neighbor) | | | | | | | | | | | | | | | | | |
| 6 | IntraPredAvailFlagA – A (Left neighbor top half) | | | | | | | | | | | | | | | | | |
| 5 | IntraPredAvailFlagE – E (Left neighbor bottom half) | | | | | | | | | | | | | | | | | |
| 4 | IntraPredAvailFlagB – B (Top neighbor) | | | | | | | | | | | | | | | | | |
| 3 | IntraPredAvailFlagC – C (Top right neighbor) | | | | | | | | | | | | | | | | | |
| 2 | IntraPredAvailFlagD – D (Top left corner neighbor) | | | | | | | | | | | | | | | | | |
| 1:0 | ChromalIntraPredMode – 2 bits to specify 1 of 4 chroma intra prediction mode, see the table in later section. | | | | | | | | | | | | | | | | | |



Inline data subfields for an Inter Macroblock

| DWord | Bit | Description | | | |
|-------|--------------|--|---|---|---|
| 4 | 31:24 | Reserved: MBZ | | | |
| | 23:16 | SubMbShapeUV (for VC-1 only) This field indicates the transformation types used for chroma components. MBZ for other codec | | | |
| | 15:8 | SubMbPredModes [bit 7:0] (Sub Macroblock Prediction Mode) This field describes the prediction mode of the sub macroblocks (four 8x8 blocks). It contains four subfields each with 2-bits, corresponding to the 4 fixed size 8x8 sub macroblocks in sequential order. This field is provided for MB with sub_mb_type equal to BP_8x8 only This field is derived from MbType for a non-BP_8x8 inter macroblock, and carries redundant information as MbType) ??? Bits [1:0]: SubMbPredMode[0] – for 8x8 Block 0 Bits [3:2]: SubMbPredMode[1] – for 8x8 Block 1 Bits [5:4]: SubMbPredMode[2] – for 8x8 Block 2 Bits [7:6]: SubMbPredMode[3] – for 8x8 Block 3 Blocks of the MB is numbered as follows : <table style="margin-left: 40px; border-collapse: collapse;"> <tr> <td style="padding-right: 20px;">0</td> <td>1</td> </tr> <tr> <td>2</td> <td>3</td> </tr> </table> Each 2-bit value [1:0] is defined as : 00 – Pred_L0 01 – Pred_L1 10 – BiPred 11 – Intra (for VC-1) ; 11 – Direct | 0 | 1 | 2 |
| 0 | 1 | | | | |
| 2 | 3 | | | | |



| DWord | Bit | Description | | | | |
|-------|-------|--|---|---|---|---|
| | 7:0 | <p>SubMbShape[bit 7:0] (Sub Macroblock Shape)</p> <p>This field describes the sub-block partitioning of each sub macroblocks (four 8x8 blocks). It contains four subfields each with 2-bits, corresponding to the 4 fixed size 8x8 sub macroblocks in sequential order.</p> <p>This field is provided for MB with sub_mb_type equal to BP_8x8 only</p> <p>This field is forced to 0 for a non-BP_8x8 inter macroblock, and effectively carries redundant information as MbType). ???</p> <p>Bits [1:0]: SubMbShape[0] – for 8x8 Block 0 Bits [3:2]: SubMbShape[1] – for 8x8 Block 1 Bits [5:4]: SubMbShape[2] – for 8x8 Block 2 Bits [7:6]: SubMbShape[3] – for 8x8 Block 3</p> <p>Blocks of the MB is numbered as follows :</p> <table style="margin-left: 40px;"> <tr> <td>0</td> <td>1</td> </tr> <tr> <td>2</td> <td>3</td> </tr> </table> <p>Each 2-bit value [1:0] is defined as :</p> <p>00 – SubMbPartWidth=8, SubMbPartHeight=8 01 – SubMbPartWidth=8, SubMbPartHeight=4 10 – SubMbPartWidth=4, SubMbPartHeight=8 11 – SubMbPartWidth=4, SubMbPartHeight=4</p> <p>For VC-1, This field indicates the transformation types used for luma components, 2 bits for each 8x8.</p> | 0 | 1 | 2 | 3 |
| 0 | 1 | | | | | |
| 2 | 3 | | | | | |
| 5 | 31:24 | <p>Frame Store ID L0[3]</p> <p>Support up to 4 Frame store ID per L0 direction, one per MB partition, if exists. See details in later section. This field specifies the frame Store ID into the Reference Picture List0 Table.</p> <p>Bit 7: Must Be One: (This is reserved for control fields in future extension, when reference index are generated instead of frame store ID) 1: indicate it is in Frame store ID format. 0: indicate it is in Reference Index format.</p> <p>Bit 6:5: reserved MBZ</p> <p>Bit 4:0 : Frame store index or Frame Store ID (Bit 4:1 is used to form the binding table index in Intel implementation)</p> | | | | |
| | 23:16 | <p>Frame Store ID L0[2]</p> <p>Support up to 4 Frame store ID per L0 direction, one per MB partition, if exists. See details in later section. This field specifies the frame Store ID into the Reference Picture List0 Table.</p> <p>Bit 7: Must Be One: (This is reserved for control fields in future extension, when reference index are generated instead of frame store ID) 1: indicate it is in Frame store ID format. 0: indicate it is in Reference Index format.</p> <p>Bit 6:5: reserved MBZ</p> <p>Bit 4:0 : Frame store index or Frame Store ID (Bit 4:1 is used to form the binding table index in Intel implementation)</p> | | | | |



| DWord | Bit | Description |
|-------|-------|--|
| | 15:8 | <p>Frame Store ID L0[1]</p> <p>Support up to 4 Frame store ID per L0 direction, one per MB partition, if exists. See details in later section. This field specifies the frame Store ID into the Reference Picture List0 Table.</p> <p>Bit 7: Must Be One: (This is reserved for control fields in future extension, when reference index are generated instead of frame store ID)</p> <p>1: indicate it is in Frame store ID format.</p> <p>0: indicate it is in Reference Index format.</p> <p>Bit 6:5: reserved MBZ</p> <p>Bit 4:0 : Frame store index or Frame Store ID (Bit 4:1 is used to form the binding table index in Intel implementation).</p> |
| | 7:0 | <p>Frame Store ID L0[0]</p> <p>Support up to 4 Frame store ID per L0 direction, one per MB partition, if exists. See details in later section. This field specifies the frame Store ID into the Reference Picture List0 Table.</p> <p>Bit 7: Must Be One: (This is reserved for control fields in future extension, when reference index are generated instead of frame store ID)</p> <p>1: indicate it is in Frame store ID format.</p> <p>0: indicate it is in Reference Index format.</p> <p>Bit 6:5: reserved MBZ</p> <p>Bit 4:0 : Frame store index or Frame Store ID (Bit 4:1 is used to form the binding table index in Intel implementation)</p> |
| 6 | 31:24 | <p>Frame Store ID L1[3]</p> <p>Support up to 4 Frame store ID per L0 direction, one per MB partition, if exists. See details in later section. This field specifies the frame Store ID into the Reference Picture List0 Table.</p> <p>Bit 7: Must Be One: (This is reserved for control fields in future extension, when reference index are generated instead of frame store ID)</p> <p>1: indicate it is in Frame store ID format.</p> <p>0: indicate it is in Reference Index format.</p> <p>Bit 6:5: reserved MBZ</p> <p>Bit 4:0 : Frame store index or Frame Store ID (Bit 4:1 is used to form the binding table index in Intel implementation)</p> |
| | 23:16 | <p>Frame Store ID L1[2]</p> <p>Support up to 4 Frame store ID per L0 direction, one per MB partition, if exists. See details in later section. This field specifies the frame Store ID into the Reference Picture List0 Table.</p> <p>Bit 7: Must Be One: (This is reserved for control fields in future extension, when reference index are generated instead of frame store ID)</p> <p>1: indicate it is in Frame store ID format.</p> <p>0: indicate it is in Reference Index format.</p> <p>Bit 6:5: reserved MBZ</p> <p>Bit 4:0 : Frame store index or Frame Store ID (Bit 4:1 is used to form the binding table index in Intel implementation)</p> |



| DWord | Bit | Description |
|-------|------|--|
| | 15:8 | <p>Frame Store ID L1[1]</p> <p>Support up to 4 Frame store ID per L0 direction, one per MB partition, if exists. See details in later section. This field specifies the frame Store ID into the Reference Picture List0 Table.</p> <p>Bit 7: Must Be One: (This is reserved for control fields in future extension, when reference index are generated instead of frame store ID)</p> <p>1: indicate it is in Frame store ID format.</p> <p>0: indicate it is in Reference Index format.</p> <p>Bit 6:5: reserved MBZ</p> <p>Bit 4:0 : Frame store index or Frame Store ID (Bit 4:1 is used to form the binding table index in Intel implementation)</p> |
| | 7:0 | <p>Frame Store ID L1[0]</p> <p>Support up to 4 Frame store ID per L0 direction, one per MB partition, if exists. See details in later section. This field specifies the frame Store ID into the Reference Picture List0 Table.</p> <p>Bit 7: Must Be One: (This is reserved for control fields in future extension, when reference index are generated instead of frame store ID)</p> <p>1: indicate it is in Frame store ID format.</p> <p>0: indicate it is in Reference Index format.</p> <p>Bit 6:5: reserved MBZ</p> <p>Bit 4:0 : Frame store index or Frame Store ID (Bit 4:1 is used to form the binding table index in Intel implementation)</p> |

3.14 Encoder StreamOut Mode Data Structure Definition

When StreamOut is enabled, per MB (and/or per Slice, per Picture) intermediated coding data (e.g. bit allocated for each MB, etc.) are sent to the memory in a fixed record format (and of fixed size) from the PAK. The per-MB records must be written in a strict raster order and with no gap (i.e. every MB regardless of its mb_type and slice type, must have an entry in the StreamOut buffer). Therefore, the consumer of the StreamOut data can offset into the StreamOut Buffer (**StreamOut Data Destination Base Address**) using individual MB addresses.

Adding per macroblock stream out for PAK is for following purposes:

- Immediate multi-pass PAK (without host or EU intervention)
 - 3200-bit conformance
 - Re-quantization
- Providing information for host for offline processing

The description for the fixed format PAK streamout record :

Streamout Pointer: Use the existing streamout pointer and enabler

Per Macroblock Information (a fixed size structure)



| Dword | Bit | Description |
|-------|--------------|--|
| 0 | 31:24 | PackedMvNum (see Decoder Stream-Out) |
| | 23 | Reserved MBZ |
| | 22-20 | MvFormat (see Decoder Stream-Out) |
| | 19:17 | CodedPatternDC (see Decoder Stream-Out) |
| | 16 | Reserved MBZ |
| | 15 | TransformFlag (see Decoder Stream-Out) |
| | 14 | MbFieldFlag (see Decoder Stream-Out) |
| | 13 | IntraMbFlag (see Decoder Stream-Out) |
| | 12:8 | MbType5Bits (see Decoder Stream-Out) |
| | 7:6 | Reserved MBZ |
| | 5:4 | IntraMbMode (see Decoder Stream-Out) |
| | 2 | SkipMbFlag (see Decoder Stream-Out) always |
| | 1:0 | InterMbMode (see Decoder Stream-Out) |
| 1 | 31:16 | MbYCnt (see Decoder Stream-Out) |
| | 15:0 | MbXCnt (see Decoder Stream-Out) |
| | 12:0 | Header Bit count (bit count due to Pre-coefficient data) per macroblock |
| 2 | 31:29 | Reserved |
| | 28:16 | MbTotBits Total Bit Count per macroblock |
| | 15:12 | Reserved |
| | 12:0 | MbHdrBits Header Bit count (bit count due to Pre-coefficient data) per macroblock |
| 3 | 31:28 AVC | Skip8x8Pattern (AVC) (see Decoder Stream-Out) |
| | 31 MPEG2 | EndSliceFlag (MPEG2) (see Decoder Stream-Out) |
| | 30 MPEG2 | NewSliceFlag (MPEG2) (see Decoder Stream-Out) |
| | 29:28 | EnableClampCoeff 1: Indicates if clamping of any coefficient of the macroblock is done for Rate Control 0: No clamping done |
| | 26 | LastMbFlag (see Decoder Stream-Out) |
| | 25 | SkipMbConvDisable (see Decoder Stream-Out) |
| | 24 | Reserved MBZ. |



| | | |
|--|-------|---|
| | 23 | MB level Inter MB conformance flag 1- Total Bit Count of a macroblock is more than Inter Conformance Max size limit in the MFX_AVC_IMG_STATE Command 0- Total Bit Count of a macroblock is not more than Inter Conformance Max size limit in the MFX_AVC_IMG_STATE Command Mutually exclusive |
| | 22 | MB level Intra MB conformance flag to trigger multi-pass 1- Total Bit Count of a macroblock is more than Intra Conformance Max size limit in the MFX_AVC_IMG_STATE Command 0- Total Bit Count of a macroblock is not more than Intra Conformance Max size limit in the MFX_AVC_IMG_STATE Command Mutually exclusive |
| | 21 | CoeffClampStat Coefficient clamp status. 1- When Coefficient clamp is applied for the MB. 0- No Coefficient clamp is applied for the MB.EnableCoeffClamp (see Decoder Stream-Out)Reserved MBZ. |
| | 20 | MB level Rate control flag 1 – when RateControlCounterEnable is set in the MFX_AVC_SLICE_STATE Command 0 - when RateControlCounterEnable is disable in the MFX_AVC_SLICE_STATE Command |
| | 19:16 | Reserved MBZ. |
| | 15:12 | MbConformDeltaQp Suggested QP delta from Conformance (QRC) for MB |
| | 11:8 | MbRateCtrlDeltaQp Suggested QP delta from RC (QRC) for MB |
| | 7 | mbz (AVC) / QScaleType (MPEG2) |
| | 6:0 | QpPrimeY (AVC) / QScaleCode (MPEG2) The luma quantization index. This is the per-MB QP value specified for the current MB. For SNB, it will always return 0xFF (unknown). |

3.14.1 PAK Multi-Pass

Multi-Pass PAK Usages:

- Intra MB 3200-bit conformance
- Inter MB Re-quantization
- Frame level Re-quantization



How to Enable Multi-Pass PAK?

- Using the existing conditional batch buffer execution capability to skip/execute the second pass
 - How to dynamically change the condition?
 - Defined one error condition register with a mask. Do HW status page update at the end of the first pass. 0 means all OK, non-zero means there is an error condition, requiring second pass. Mask is used by the host to control what kind of multi-pass is intended.
 - For example, one error bit is 3200-bit conformance violation. Another error bit is the total bit count exceeds (too much or too little) the target range (need to define the target range in the state).
 - **The logic perfectly fits in the conditional batch buffer control logic that VCS has today in GT. There is no additional logic need to be added in VCS to support media functionality. (Batch Buffer Skip: This field only takes effect if Compare Semaphore is set and the value at Semaphore Address is NOT greater than the Semaphore Data Dword).**
- Adding a picture level state command to enable and control the behavior of the second pass PAK
 - How to control the re-PAK? Added 3 conformance flags (error registers) in the per-MB streamout. Then the error control is based on the error register and the mask defined in picture level states. There are 8 register flags defined out of which only the 3200-bit case has usage model defined for today. The rest are left for future usage.

Issues and Limitations:

- There is no programmable engine in MFX for flexible control: Therefore, whatever we have defined must consider flexibility

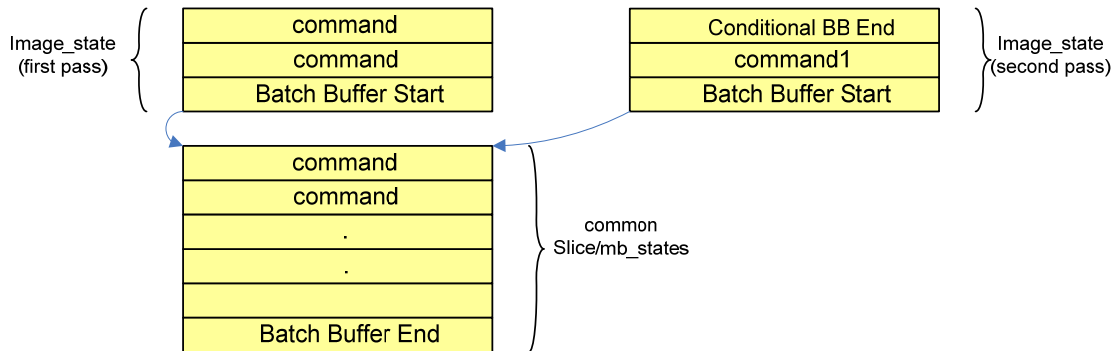
Following 2 MI packets are used inside VCS without any change to support Multipass-PAK behaviour.

- MI_Conditional_Batch_Buffer_End
- Memory Interface Registers

3.14.2 Driver Usage

Driver places Image states in one batch buffer and all slice level and macroblock level states into another batch buffer and link 2 batch buffers. Also replicate Image states with multipass changes in another batch buffer link them to slice/macroblock batch buffer. In this way, only Image states are replicated but not the slice/macroblock states. The image states includes all buffers defined at image(indirectMV, original pixel buffer, etc). Following changes are needed in the Multipass Image State,

- **Reset- Stream-Out Enable(disable stream out in the second pass)**
- **Set- MacroblockStatEnable (enable reading of macroblock status buffer)**
- **Reset- 3200-bit conformance (do not report 3200-bit conformance)**



Define Conditional Batch Buffer End for CS/VCSVINunit

3.15 Monochrome Picture Processing

Monochrome picture is specified using the Surface State with Surface Format of 12. Therefore, MFX hardware, in either decode or encode mode, doesn't generation any read or write traffic for U/V components. *Motivation for this bandwidth optimization is that monochrome video coding might be used for wireless display.*

For Encoder :

- 1) No read in UV original components.
- 2) processing UV component - no
- 3) reconstructed UV component reference picture - no
- 4) filter UV component - no

For Decoder :

- 1) VLD mode : no color component coming out of core in Monochrome mode and so no processing and not writing output
- 2) IT mode : there is no color component in the coefficient buffer, and so no processing and not writing output

3.16 Context Switch

There is no pre-emption for the BCS pipeline; hence every command buffer is required to contain all the states setup (preamble). Specifically, CPU can not interrupt the BCS-BSD pipe, to stop the operation in the middle of decoding a bitstream data.

Switch of contexts can only be performed at picture boundary.

No state need to be saved



3.17 Pipeline Flush

Implicit flush for AVC and VC1 is performed at the end of Slice : for MPEG2 is done when a new image/picture command is issued. Because MPEG2 a slice can be one MB, no point to flush. MPEG2 will snoop the next command if it is an img_state command.

Explicit flush MI (1 bit to do media pipeline vs Gx pipeline) flush and cache flush (switch reference frame) – MI flush has bit to do cache flush. MI flush is for driver synchronization. MMIO Interface

A set of registers are defined and accessible through MMIO interface to serve multiple purposes :

- Use for system configuration
- For accessing Performance counters
- WatchDog Timer – VCS command movement

3.18 AVC Error Handling and Error Concealment

Design objective is to keep the playback continue going forward in time even when errors are encountered (detected or not). Hence, the system will reconstruct and display the new picture as much as possible, without dropping the current picture or repeating the previous decoded picture. Another aspect of the design is to minimize the interaction/intervention of driver software when non-fatal error is encountered. For example, the driver can always insert a Phantom Slice at the end of a picture, regardless of whether the picture contains error or not.

Error handling is performed only at the Slice boundary, and in between pictures.

To catch the case when the VC1_BSD has finished decoding the current picture but the last macroblock is not reached, a Phantom Slice is required to conceal such error condition. Therefore, it is required to submit an additional VC1_BSD_OBJECT command for the Phantom Slice at the end of each picture. A Phantom Slice is one that has the **Indirect Data Length** set to 0 and **SliceStartVertPos** set to the height of the current picture in macroblocks, indicating that it is just outside the picture boundary.

Note that due to the nature of splitting the decoding pipeline into the front end BSD and the back end picture reconstruction, error concealment must be performed by the front end BSD. Otherwise, the commands generated for the back end picture reconstruction may be incomplete and therefore cause further problems (including hardware hang conditions).



3.18.1 AVC Design Assumptions

Operating Assumptions:

- The BSD unit is either actively decoding a picture, or in between pictures when errors occur.
- At the beginning of decoding a new picture (the first slice of a picture), the current MB location registers are reset to 0 by the hardware (??? Automatically or through State Command)
- The end of a picture is tracked by comparing the current MB location (just being decoded) to the the picture size parameters (in MB unit) in the BSD unit. The BSD unit will never attempt to decode beyond the picture size programmed, even when there are data remained at the BSD input.
- The end of a picture sequence is only known in the driver. BSD unit will continue operating until no more decoding request commands are injected or when there is an optional time-out by the driver.
- BSD unit has no notion of picture structure, it only deals with Slice as instructed and directed by the State Commands. It does not know it is decoding the last slice of a picture.
- There is no hardware Start_Code_Detector at the front end of the BSD unit, so the end of Slice signal can only be detected and generated internally.

3.18.2 AVC Error Concealment and Recovery Strategies

Depending on the category and position of error being detected, a concealment strategy will be adopted.

- Picture Repeat
 - Handled by the display controller – when it does not receive a new picture to display at the appropriate time period
- Hardware MB Fillings (as described below)
- S/W Re-sync at the next Slice or picture boundary
 - Dropping a Slice, or
 - The driver can ignore the rest of the bitstream until the start of the next NAL is detected. That is to start decoding again at slice boundary
 - dropping a picture entirely.
 - The driver can decide to ignore all subsequent slices until a new IDR has arrived, and to start decoding again from there on.



3.18.3 AVC Error Handling

3.18.3.1 S/W Driver

Provide high level error handling capabilities – interrupts and/or polling. Perform error handling and recovery in between Slices and Pictures boundaries.

It also checks for the bitstream syntax errors at and above Slice Header Layers, before proceeding with the Slice decoding.

Currently there is no communication of error handling at the application level.

There are also two error status bits generated in the hardware that the S/W can poll regularly :

- Force_In_Progress
 - A hardware Status Bit (read-only) to indicate filling missing MBs is in progress
- Error_In_Slice
 - A hardware Status Bit (read-only) to indicate an error has detected in the HW, including VLC Decoding Error, BSD Decoded Data Error and MPR Error

3.18.3.2 Hardware MB Filling

This is the only error handling mechanism built-in the BSD unit. It will automatically perform error detection and error concealment in the Slice Data Layer at the Slice boundary. There is no direct S/W control.

When decoding a Slice (contains multiple MBs), internally the BSD unit will keep track of the number of MBs being decoded so far. This number also represents the position of the corresponding MB in a picture, assuming a raster-scan MB order (top to bottom, left to right). Hence, the Current_MB_Number register is maintained and updated in the hardware automatically.

For each Slice Data, there is also a Start_MB_Number presented at the corresponding Slice Header. This number specifies the picture location of the very first MB decoded from this Slice Data.

Just before start decoding the current new Slice Data, the last value of the Current_MB_Number register (holding the last MB position of the previously decoded Slice) is compared against the Start_MB_Number. If there is a match, the decoding of current new Slice continues. Otherwise, error concealment is immediately triggered.

There are 2 erroneous cases :

- Current_MB_Number > Start_MB_Number
 - Current_MB_Number will force to equal to Start_MB_Number, and decoding continues at the newly-adjusted Current_MB_Number location.
 - If Current_MB_Number has reached the end of the picture, any data remains in the input must be discarded.
- Current_MB_Number < Start_MB_Number



- The gap is automatically filled, so that the decoding can continue at the after-filled Current_MB_Number location
 - The starting position for MB filling is adjustable, and is equal to $\text{Current_MB_Number} - \text{MIN}\{\text{Decoded_MB_Count_Prev}, [(\text{MBAFF} ? 2:1) * \text{Force_Skip_Rewind}]\}$
- The MBs used to fill the gap are derived as follows :
 - Hardware does not support Intra Concealment Mode
 - Inter Concealment Mode
 - The missing MBs are concealed by copying co-located MBs from the frame store specified as the error concealment frame store during the slice start command. By default the top of the list_0 is used.
- If the current premature completed slice is the last slice of a picture, SW will need to inject a Phantom Slice to cause the HW to fill the missing MBs to the end of the picture
 - A Phantom Slice with Start_MB_NUMBER set beyond or equal to the picture size
 - Decoding stops when the picture size has reached.

For example: a picture contains 3 Slices

- Slice 1 – contains 100 MBs
- Slice 2 – contains 200 MBs
- Slice 3 – contains 50 MBs

After decoded the Slice Data of Slice 1, the Current_MB_Number is rest at 80, When start to decode Slice 2, its Start_MB_Number is set at 100. Hence, the hardware will automatically fill-in the missing MBs from 81 to 99. Then Slice 2 will continue its decoding from location 100.

Erroneous MB decoding will never pass to the next picture. Since, internally the BSD unit has stored the size of the picture (in MB unit), and it will stop further processing (error fillings, or MB decoding) once this boundary value has reached.

Refer to later section on supporting multiple stream decoding, in which the Current_MB_Number register can be programmed from the driver to allow additional flexibility in error handling.

Overflow impact ILDB parameters in BSD – corrupted the row store, live with it and continue but need to test not to hang the machine.

3.18.3.3 Error Statistics

TBD



3.18.3.4 WatchDog Timers

At the driver level an additional error detection mechanism is needed. In the situation where errors have stalled the MFD Engine from progressing for a predefined period of time, there must be a way to tell that the hardware is hung. Watch-dog timers are preset with some known timing constraints that are related to the decoding, e.g. an expected decoding time of a slice or a picture. Once the threshold has reached, a stall condition is assumed, and action will take place to reset the MFD Engine, and up to the driver and application to decide where to start decoding again or simply abort the decoding processing.

3.19 VC1 Error Handling and Error Concealment

Same approach as AVC, except that VC1 also has the additional information for end of Slice detection. VC1 can snoop the starting MB address of next slice, so when processing the current slice, the VC1-BSD engine already knows the last MB address.

4. Reset internal MB address on MB address overflow - VC1 is not the same as AVC, as VC1 has the last MB address, so it will stop right at it all the time.
 - b. VC1 will flush out the extra data in the bitstream for the current slice by VIN
5. Premature ending of a slice
 - d. no rewind mechanism in VC1
 - e. compare the end of MB address to the start address of next slice
 - 4) always complete the current erroneous MB to the end even if there is error, and the error concealment will start in the next MB address, only if the error concealment is enabled; otherwise stop there, and send out end of slice, and start decoding the next slice – so leaving a hole in the picture.
 1. need to double check if error concealment is disabled, will the hole causing problem in the VFE pipe ????
 - f. Error concealment regardless of picture type - interlace, interlace frame or progressive frame
 - 5) Intra
 1. I-picture
 2. Only DC is predicted from neighbor, supply one DC coefficient per block and send out EOB. That is all AC coefficients are effectively set to 0.
 3. DC coefficient is sent to IT for further processing
 - 6) Inter
 4. assume to be P-MB in both P and B-Slice
 5. 1MV Skip MB



6. Default to frame MB and use frame reference motion for interlaced frame picture
 7. No IT, only MC ($MV = 0,0$, use the nearest reference frame)
6. There is no phantom slice being sent from driver at the end of each VC1 picture

3.20 MPEG2 Error Handling and Error Concealment

No Error Handling and Error Concealment defined in MPEG2

3.21 Concurrent, Multiple Video Stream Decoding Support

The natural place for switching across multiple streams is at the Slice boundary. Each Slice is a self-sustained unit of compressed video data and has no dependency with its neighbors (except for the Deblocking process). In addition, there is no interruptability within a Slice. However, when ILDB is invoked, the processing of some MBs will require neighbour MB information that cross the Slice boundary. Hence, to limit the buffering requirement, in this version of hardware design, the stream switching can only be performed at the picture boundary instead.

Frame boundary.