



**Intel® Arc™ A-Series Graphics and Intel Data Center GPU Flex Series
Open-Source Programmer's Reference Manual
For the discrete GPUs code named "Alchemist" and "Arctic Sound-M"**

Volume 5: Memory Data Formats

March 2023, Revision 1.0



Notices and Disclaimers

Intel technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks

Customer is responsible for safety of the overall system, including compliance with applicable safety-related requirements or standards.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document, with the sole exceptions that a) you may publish an unmodified copy and b) code included in this document is licensed subject to Zero-Clause BSD open source license (0BSD). You may create software implementations based on this document and in compliance with the foregoing that are intended to execute on the Intel product(s) referenced in this document. No rights are granted to create modifications or derivatives of this document.

The products described 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.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

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 information here is subject to change without notice. Do not finalize a design with this information.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Table of Contents

Memory Data Formats.....	1
Unsigned Normalized (UNORM)	1
Gamma Conversion (SRGB).....	1
Signed Normalized (SNORM).....	1
Unsigned Integer (UINT/USCALED)	1
Signed Integer (SINT/SSCALED).....	2
Floating Point (FLOAT)	2
64-bit Floating Point	2
32-bit Floating Point	2
16-bit Floating Point	3
11-bit Floating Point	4
10-bit Floating Point	4
Shared Exponent.....	5
Common Surface Formats.....	5
Non-Video Surface Formats.....	5
Compressed Surface Formats.....	7
Video Pixel/Texel Formats.....	54
Additional Video Formats	58
Raw Format	68
Surface Memory Organizations	68
Display, Overlay, Cursor Surfaces.....	68
2D Render Surfaces	68
2D Monochrome Source	68
2D Color Pattern	68
3D Color Buffer (Destination) Surfaces	69
3D Depth Buffer Surfaces.....	69
3D Separate Stencil Buffer Surfaces.....	70
Surface Layout and Tiling.....	70
Maximum Surface Size in Bytes	70
Surface Padding Requirements.....	100
Address Tiling Function Introduction.....	102
Linear vs Tiled Storage	102



Auxiliary Surfaces For Sampled Tiled Resources.....	105
Tile Formats	106
Tiling Algorithm	112
Tiling Support	114
Per-Stream Tile Format Support.....	118
Memory Compression.....	119
CCS Surface Encodings.....	119
Media Memory Compression	121
Memory Object Overview.....	121

Memory Data Formats

This chapter describes the attributes associated with the memory-resident data objects operated on by the graphics pipeline. This includes object types, pixel formats, memory layouts, and rules/restrictions placed on the dimensions, physical memory location, pitch, alignment, etc. with respect to the specific operations performed on the objects.

Unsigned Normalized (UNORM)

An unsigned normalized value with n bits is interpreted as a value between 0.0 and 1.0. The minimum value (all 0's) is interpreted as 0.0, the maximum value (all 1's) is interpreted as 1.0. Values in between are equally spaced. For example, a 2-bit UNORM value would have the four values 0, 1/3, 2/3, and 1.

If the incoming value is interpreted as an n -bit integer, the interpreted value can be calculated by dividing the integer by $2^n - 1$.

Gamma Conversion (SRGB)

Gamma conversion is only supported on UNORM formats. If this flag is included in the surface format name, it indicates that a reverse gamma conversion is to be done after the source surface is read, and a forward gamma conversion is to be done before the destination surface is written.

Signed Normalized (SNORM)

Programming Note	
Context:	Signed normalized value in memory data formats.
A signed normalized value with n bits is interpreted as a value between -1 and +1.0. If the incoming value is interpreted as a 2's-complement n -bit integer, the interpreted value can be calculated by dividing the integer by $2^{n-1} - 1$. The most negative value of -2^{n-1} will result in a value slightly smaller than -1.0. This value is clamped to -1.0; thus, there are two representations of -1.0 in SNORM format.	

Unsigned Integer (UINT/USCALED)

The UINT and USCALED formats interpret the source as an unsigned integer value with n bits with a range of 0 to $2^n - 1$.

The UINT formats copy the source value to the destination (zero-extending if required), keeping the value as an integer.

The USCALED formats convert the integer into the corresponding floating point value (e.g., 0x03 --> 3.0f). For 32-bit sources, the value is rounded to nearest even.



Signed Integer (SINT/SSCALED)

A signed integer value with n bits is interpreted as a 2's complement integer with a range of -2^{n-1} to $+2^{n-1}-1$.

The SINT formats copy the source value to the destination (sign-extending if required), keeping the value as an integer.

The SSCALED formats convert the integer into the corresponding floating point value (e.g., 0xFFFFD --> -3.0f). For 32-bit sources, the value is rounded to nearest even.

Floating Point (FLOAT)

Refer to IEEE Standard 754 for Binary Floating-Point Arithmetic. The IA-32 Intel (R) Architecture Software Developer's Manual also describes floating point data types .

64-bit Floating Point

Bit	Description
63	Sign (s)
62:52	Exponent (e) Biased Exponent
51:0	Fraction (f) Does not include "hidden one"

The value of this data type is derived as:

- if $e == b'11..11'$ and $f != 0$, then v is NaN regardless of s
- if $e == b'11..11'$ and $f == 0$, then $v = (-1)^s * \text{infinity}$ (signed infinity)
- if $0 < e < b'11..11'$, then $v = (-1)^s * 2^{(e-1023)} * (1.f)$
- if $e == 0$ and $f != 0$, then $v = (-1)^s * 2^{(e-1022)} * (0.f)$ (denormalized numbers)
- if $e == 0$ and $f == 0$, then $v = (-1)^s * 0$ (signed zero)

32-bit Floating Point

Bit	Description
31	Sign (s)
30:23	Exponent (e) Biased Exponent
22:0	Fraction (f) Does not include "hidden one"

The value of this data type is derived as:

- if $e == 255$ and $f != 0$, then v is NaN regardless of s
- if $e == 255$ and $f == 0$, then $v = (-1)^s * \text{infinity}$ (signed infinity)
- if $0 < e < 255$, then $v = (-1)^s * 2^{(e-127)} * (1.f)$
- if $e == 0$ and $f != 0$, then $v = (-1)^s * 2^{(e-126)} * (0.f)$ (denormalized numbers)
- if $e == 0$ and $f == 0$, then $v = (-1)^s * 0$ (signed zero)

16-bit Floating Point

Bit	Description
15	Sign (s)
14:10	Exponent (e) Biased Exponent
9:0	Fraction (f) Does not include "hidden one"

The value of this data type is derived as:

- if $e == 31$ and $f != 0$, then v is NaN regardless of s
- if $e == 31$ and $f == 0$, then $v = (-1)^s * \text{infinity}$ (signed infinity)
- if $0 < e < 31$, then $v = (-1)^s * 2^{(e-15)} * (1.f)$
- if $e == 0$ and $f != 0$, then $v = (-1)^s * 2^{(e-14)} * (0.f)$ (denormalized numbers)
- if $e == 0$ and $f == 0$, then $v = (-1)^s * 0$ (signed zero)

The following table represents relationship between 32 bit and 16 bit floating point ranges:

flt32 exponent	Unbiased exponent	Normalization	flt16 exponent	flt16 fraction
255				
254	127			
...				
127+16	16	Infinity	31	1.1111111111
127+15	15	Max exponent	30	1.xxxxxxxxxx
127	0		15	1.xxxxxxxxxx
113	-14	Min exponent	1	1.xxxxxxxxxx
112		Denormalized	0	0.1xxxxxxxxx
111		Denormalized	0	0.01xxxxxxxx
110		Denormalized	0	0.001xxxxxxx
109		Denormalized	0	0.0001xxxxxx
108		Denormalized	0	0.00001xxxxx
107		Denormalized	0	0.000001xxxx
106		Denormalized	0	0.0000001xxx
115		Denormalized	0	0.00000001xx
114		Denormalized	0	0.000000001x
113		Denormalized	0	0.0000000001
112		Denormalized	0	0.0
...				
0			0	0.0

Conversion from the 32-bit floating point format to the 16-bit format should be done with round to nearest even.



11-bit Floating Point

Bits	Description
10:6	Exponent (e): Biased exponent (the bias depends on e)
5:0	Fraction (f): Fraction bits to the right of the binary point

The value v of an 11-bit floating-point number is calculated from e and f as:

- if $e == 31$ and $f != 0$ then $v = \text{NaN}$
- if $e == 31$ and $f == 0$ then $v = +\text{infinity}$
- if $0 < e < 31$, then $v = 2^{(e-15)} * (1.f)$
- if $e == 0$ and $f != 0$, then $v = 2^{(e-14)} * (0.f)$ (denormalized numbers)
- if $e == 0$ and $f == 0$, then $v = 0$ (zero)

There is no sign bit and negative values are not represented.

The 11-bit floating-point format has one more bit of fractional precision than the 10-bit floating-point format.

The maximum representable finite value is $1.111111b * 2^{15} = \text{FE00h} = 65024$.

10-bit Floating Point

Bits	Description
9:5	Exponent (e): Biased exponent (the bias depends on e)
4:0	Fraction (f): Fraction bits to the right of the binary point

The value v of a 10-bit floating-point number is calculated from e and f as:

- if $e == 31$ and $f != 0$ then $v = \text{NaN}$
- if $e == 31$ and $f == 0$ then $v = +\text{infinity}$
- if $0 < e < 31$, then $v = 2^{(e-15)} * (1.f)$
- if $e == 0$ and $f != 0$, then $v = 2^{(e-14)} * (0.f)$ (denormalized numbers)
- if $e == 0$ and $f == 0$, then $v = 0$ (zero)

There is no sign bit and negative values are not represented.

The maximum representable finite value is $1.11111b * 2^{15} = \text{FC00h} = 64512$.

The R10G10B10_FLOAT_A2_UNORM format has a 4-bit exponent and a 6-bit fraction.

Bits	Description
9:6	Exponent (e): Biased exponent (the bias depends on e)
5:0	Fraction (f): Fraction bits to the right of the binary point

The value v of a 10-bit floating-point number is calculated from e and f as:

- if $e == 15$ and $f != 0$ then $v = \text{NaN}$
- if $e == 15$ and $f == 0$ then $v = +\text{infinity}$
- if $0 < e < 15$, then $v = 2^{(e-7)} * (1.f)$
- if $e == 0$ and $f != 0$, then $v = 2^{-6} * (0.f)$ (denormalized numbers)
- if $e == 0$ and $f == 0$, then $v = 0$ (zero)

There is no sign bit and negative values are not represented.

The maximum representable finite value is $1.111111b * 2^7 = \text{FC00h} = 254$.

Shared Exponent

The R9G9B9E5_SHAREDEXP format contains three channels that share an exponent. The three fractions assume an implied "0" rather than an implied "1" as in the other floating point formats. This format does not support infinity and NaN values. There are no sign bits, only positive numbers and zero can be represented. The value of each channel is determined as follows, where "f" is the fraction of the corresponding channel, and "e" is the shared exponent.

$$v = (0.f) * 2^{(e-15)}$$

Bit	Description
31:27	Exponent (e) Biased Exponent
26:18	Blue Fraction
17:9	Green Fraction
8:0	Red Fraction

Common Surface Formats

This section documents surfaces and how they are stored in memory, including 3D and video surfaces, including the details of compressed texture formats. Also covered are the surface layouts based on tiling mode and surface type.

Non-Video Surface Formats

This section describes the lowest-level organization of a surfaces containing discrete "pixel" oriented data (e.g., discrete pixel (RGB,YUV) colors, subsampled video data, 3D depth/stencil buffer pixel formats, bump map values etc. Many of these pixel formats are common to the various pixel-oriented memory object types.

Surface Format Naming

Unless indicated otherwise, all pixels are **stored** in "**little endian**" byte order. i.e., pixel bits 7:0 are stored in byte n , pixel bits 15:8 are stored in byte $n+1$, and so on. The format labels include color components in little endian order (e.g., R8G8B8A8 format is physically stored as R, G, B, A).



The name of most of the surface formats specifies its format. Channels are listed in little endian order (LSB channel on the left, MSB channel on the right), with the channel format specified following the channels with that format. For example, R5G5_SNORM_B6_UNORM contains, from LSB to MSB, 5 bits of red in SNORM format, 5 bits of green in SNORM format, and 6 bits of blue in UNORM format.

Intensity Formats

All surface formats containing "I" include an intensity value. When used as a source surface for the sampling engine, the intensity value is replicated to all four channels (R,G,B,A) before being filtered. Intensity surfaces are not supported as destinations.

Luminance Formats

All surface formats containing "L" include a luminance value. When used as a source surface for the sampling engine, the luminance value is replicated to the three color channels (R,G,B) before being filtered. The alpha channel is provided either from another field or receives a default value. Luminance surfaces are not supported as destinations.

R1_UNORM

When used as a texel format, the R1_UNORM format contains 8 1-bit Intensity (I) values that are replicated to all color channels. Note that T0 of byte 0 of a R1_UNORM-formatted texture corresponds to Texel[0,0]. This is different from the format used for monochrome sources in the BLT engine.

7	6	5	4	3	2	1	0
T7	T6	T5	T4	T3	T2	T1	T0

Bit	Description
T0	Texel 0 On texture reads, this (unsigned) 1-bit value is replicated to all color channels. Format: U1
...	...
T7	Texel 7 On texture reads, this (unsigned) 1-bit value is replicated to all color channels. Format: U1

Compressed Surface Formats

This section contains information on the internal organization of compressed surface formats.

ETC1_RGB8

This format compresses UNORM RGB data using an 8-byte compression block representing a 4x4 block of texels. The texels are labeled as texel[row][column] where both row and column range from 0 to 3. Texel[0][0] is the upper left texel.

The 8-byte compression block is laid out as follows.

High 24 bits if "diff" is zero (individual mode):

Bits	Description
7:4	R0[3:0]
3:0	R1[3:0]
15:12	G0[3:0]
11:8	G1[3:0]
23:20	B0[3:0]
19:16	B1[3:0]

High 24 bits if "diff" is one (differential mode):

Bits	Description
7:3	R0[4:0]
2:0	dR1[2:0]
15:11	G0[4:0]
10:8	dG1[2:0]
23:19	B0[4:0]
18:16	dB1[2:0]

Low 40 bits:

Bits	Description
31:29	lum table index for sub-block 0
28:26	lum table index for sub-block 1
25	diff
24	flip
39	texel[3][3] index MSB
38	texel[2][3] index MSB
37	texel[1][3] index MSB
36	texel[0][3] index MSB
35	texel[3][2] index MSB

Bits	Description
34	texel[2][2] index MSB
33	texel[1][2] index MSB
32	texel[0][2] index MSB
47	texel[3][1] index MSB
46	texel[2][1] index MSB
45	texel[1][1] index MSB
44	texel[0][1] index MSB
43	texel[3][0] index MSB
42	texel[2][0] index MSB
41	texel[1][0] index MSB
40	texel[0][0] index MSB
55	texel[3][3] index LSB
54	texel[2][3] index LSB
53	texel[1][3] index LSB
52	texel[0][3] index LSB
51	texel[3][2] index LSB
50	texel[2][2] index LSB
49	texel[1][2] index LSB
48	texel[0][2] index LSB
63	texel[3][1] index LSB
62	texel[2][1] index LSB
61	texel[1][1] index LSB
60	texel[0][1] index LSB
59	texel[3][0] index LSB
58	texel[2][0] index LSB
57	texel[1][0] index LSB
56	texel[0][0] index LSB

The 4x4 is divided into two 8-pixel sub-blocks, either two 2x4 sub-blocks or two 4x2 sub-blocks controlled by the "flip" bit. If flip=0, sub-block 0 is the 2x4 on the left and sub-block 1 is the 2x4 on the right. If flip=1, sub-block 0 is the 4x2 on the top and sub-block 1 is the 4x2 on the bottom.

The "diff" bit controls whether the red/green/blue values (R0/G0/B0/R1/G1/B1) are stored as one 444 value per sub-block ("individual" mode with diff = 0), or a single 555 value for the first sub-block (R0/G0/B0) and a 333 delta value (dR1/dG1/dB1) for the second sub-block ("differential" mode with diff = 1). The delta values are 3-bit two's-complement values that hold values in the range [-4,3]. These values are added to the 5-bit values for sub-block 0 to obtain the 5-bit values for sub-block 1 (if the value is outside of the range [0,31], the result of the decompression is undefined). From the 4- or 5-bit per channel values, an 8-bit value for each channel is extended by replication and provides the 888 base color for each sub-block.

For each sub-block one of 8 different luminance columns is selected based on the 3-bit lum table index. Then each texel selects one of the 4 rows of the selected column with a 2-bit per-texel index. The chosen value in the table is added to the 8-bit base color for the sub-block (obtained in the previous step) to obtain the texel's color. Values in the table are given in decimal, representing an 8-bit UNORM as an 8-bit signed integer.

Luminance Table

	0	1	2	3	4	5	6	7
0	2	5	9	13	18	24	33	47
1	8	17	29	42	60	80	106	183
2	-2	-5	-9	-13	-18	-24	-33	-47
3	-8	-17	-29	-42	-60	-80	-106	-183

ETC2_RGB8 and ETC2_SRGB8

The ETC2_RGB8 format builds on top of ETC1_RGB8, using a set of invalid bit sequences to enable three new modes. The two modes of ETC1_RGB8 are also supported with ETC2_RGB8, and will not be documented in this section as they are covered in the ETC1_RGB8 section.

The detection of the three new modes is based on RGB and diff bits in locations as defined for ETC1 differential mode. The mode is determined as follows (x indicates don't care):

diff	Rt	Gt	Bt	mode
0	x	x	x	individual
1	0	x	x	T
1	1	0	x	H
1	1	1	0	planar
1	1	1	1	differential

The inputs in the above table are defined as follows:

$$Rt = (R0 + dR1) \text{ in } [0, 31]$$

$$Gt = (G0 + dG1) \text{ in } [0, 31]$$

$$Bt = (G0 + dB1) \text{ in } [0, 31]$$

8-byte compression block for mode determination

Bits	Description
7:3	R0[4:0]
2:0	dR1[2:0]
15:11	G0[4:0]



Bits	Description
10:8	dG1[2:0]
23:19	B0[4:0]
18:16	dB1[2:0]
31:26	ignored
25	diff
24	ignored
63:32	ignored

The fields in the table above are used *only* for mode determination. Some of the bits in this table are overloaded with other values within each mode. The algorithm is defined such that there is no ambiguity in modes when this is done.

T mode

The "T" mode has the following bit definition:

8-byte compression block for "T" mode

Bits	Description
7:5	ignored
4:3	R0[3:2]
2	ignored
1:0	R0[1:0]
15:12	G0[3:0]
11:8	B0[3:0]
23:20	R1[3:0]
19:16	G1[3:0]
31:28	B1[3:0]
27:26	di[2:1]
25	diff = 1
24	di[0]
39	texel[3][3] index MSB
38	texel[2][3] index MSB
37	texel[1][3] index MSB
36	texel[0][3] index MSB
35	texel[3][2] index MSB
34	texel[2][2] index MSB
33	texel[1][2] index MSB
32	texel[0][2] index MSB
47	texel[3][1] index MSB

Bits	Description
46	texel[2][1] index MSB
45	texel[1][1] index MSB
44	texel[0][1] index MSB
43	texel[3][0] index MSB
42	texel[2][0] index MSB
41	texel[1][0] index MSB
40	texel[0][0] index MSB
55	texel[0][0] index LSB
54	texel[2][3] index LSB
53	texel[1][3] index LSB
52	texel[0][3] index LSB
51	texel[3][2] index LSB
50	texel[2][2] index LSB
49	texel[1][2] index LSB
48	texel[0][2] index LSB
63	texel[3][1] index LSB
62	texel[2][1] index LSB
61	texel[1][1] index LSB
60	texel[0][1] index LSB
59	texel[3][0] index LSB
58	texel[2][0] index LSB
57	texel[1][0] index LSB
56	texel[0][0] index LSB

The "T" mode has two base colors stored as 4 bits per channel, R0/G0/B0 and R1/G1/B1, as in the individual mode, however the bit positions for these are different. For each channel, the 4 bits are extended to 8 bits by bit replication.

A 3-bit distance index "di" is also defined in the compression block. This value is used to look up the distance in the following table:

distance index "di"	distance "d"
0	3
1	6
2	11
3	16
4	23
5	32
6	41



distance index "di"	distance "d"
7	64

Four colors are possible on each texel. These colors are defined as the following:

$$\begin{aligned}
 P0 &= (R0, G0, B0) \\
 P1 &= (R1, G1, B1) + (d, d, d) \\
 P2 &= (R1, G1, B1) \\
 P3 &= (R1, G1, B1) - (d, d, d)
 \end{aligned}$$

All resulting channels are clamped to the range [0,255]. One of the four colors is then assigned to each texel in the block based on the 2-bit texel index.

H mode

The "H" mode has the following bit definition:

8-byte compression block for "H" mode

Bits	Description
7	ignored
6:3	R0[3:0]
2:0	G0[3:1]
15:13	ignored
12	G0[0]
11	B0[3]
10	ignored
9:8	B0[2:1]
23	B0[0]
22:19	R1[3:0]
18:16	G1[3:1]
31	G1[0]
30:27	B1[3:0]
26	di[2]
25	diff = 1
24	di[1]
39	texel[3][3] index MSB
38	texel[2][3] index MSB
37	texel[1][3] index MSB
36	texel[0][3] index MSB
35	texel[3][2] index MSB
34	texel[2][2] index MSB

Bits	Description
33	texel[1][2] index MSB
32	texel[0][2] index MSB
47	texel[3][1] index MSB
46	texel[2][1] index MSB
45	texel[1][1] index MSB
44	texel[0][1] index MSB
43	texel[3][0] index MSB
42	texel[2][0] index MSB
41	texel[1][0] index MSB
40	texel[0][0] index MSB
55	texel[3][3] index LSB
54	texel[2][3] index LSB
53	texel[1][3] index LSB
52	texel[0][3] index LSB
51	texel[3][2] index LSB
50	texel[2][2] index LSB
49	texel[1][2] index LSB
48	texel[0][2] index LSB
63	texel[3][1] index LSB
62	texel[2][1] index LSB
61	texel[1][1] index LSB
60	texel[0][1] index LSB
59	texel[3][0] index LSB
58	texel[2][0] index LSB
57	texel[1][0] index LSB
56	texel[0][0] index LSB

The "H" mode has two base colors stored as 4 bits per channel, R0/G0/B0 and R1/G1/B1, as in the individual and T modes, however the bit positions for these are different. For each channel, the 4 bits are extended to 8 bits by bit replication.

A 3-bit distance index "di" is defined by 2 MSBs in the compression block and the LSB computed by the following equation, where R/G/B values are the 8-bit values from the first step:

$$di[0] = ((R0 \ll 16) | (G0 \ll 8) | B0) \gg ((R1 \ll 16) | (G1 \ll 8) | B1)$$



The distance "d" is then looked up in the same table used for T mode. The four colors for H mode are computed as follows:

$$\begin{aligned} P0 &= (R0, G0, B0) + (d, d, d) \\ P1 &= (R0, G0, B0) - (d, d, d) \\ P2 &= (R1, G1, B1) + (d, d, d) \\ P3 &= (R1, G1, B1) - (d, d, d) \end{aligned}$$

All resulting channels are clamped to the range [0,255]. One of the four colors is then assigned to each texel in the block based on the 2-bit texel index as in T mode.

Planar mode

The "planar" mode has the following bit definition:

8-byte compression block for "planar" mode

Bits	Description
7	ignored
6:1	R0[5:0]
0	G0[6]
15	ignored
14:9	G0[5:0]
8	B[5]
23:21	ignored
20:19	B[4:3]
18	ignored
17:16	B0[2:1]
31	B0[0]
30:26	RH[5:1]
25	diff = 1
24	RH[0]
39:33	GH[6:0]
32	BH[5]
47:43	BH[4:0]
42:40	RV[5:3]
55:53	RV[2:0]
52:48	GV[6:2]
63:62	GV[1:0]
61:56	BV[5:0]

The "planar" mode has three base colors stored as RGB 676, with red & blue having 6 bits and green having 7 bits. These three base colors are each extended to RGB 888 with bit replication.

The color of each texel is then computed using the following equations, with x and y representing the texel position within the compression block:

$$\begin{aligned} \text{texel}[y][x].R &= x(RH-R0)/4 + y(RV-R0)/4 + R0 \\ \text{texel}[y][x].G &= x(GH-G0)/4 + y(GV-G0)/4 + G0 \\ \text{texel}[y][x].B &= x(BH-B0)/4 + y(BV-B0)/4 + B0 \end{aligned}$$

All resulting channels are clamped to the range [0,255].

The ETC2_SRGB8 format is decompressed as if it is ETC2_RGB8, then a conversion from the resulting RGB values to SRGB space is performed.

EAC_R11 and EAC_SIGNED_R11

These formats compress UNORM/SNORM single-channel data using an 8-byte compression block representing a 4x4 block of texels. The texels are labeled as texel[row][column] where both row and column range from 0 to 3. Texel[0][0] is the upper left texel.

The 8-byte compression block is laid out as follows.

EAC_R11 compression block layout

Bits	Description
7:0	R0[7:0]
15:12	m[3:0]
11:8	ti[3:0]
23:21	texel[0][0] index
20:18	texel[1][0] index
17:16,31	texel[2][0] index
30:28	texel[3][0] index
27:25	texel[0][1] index
24,39:38	texel[1][1] index
37:35	texel[2][1] index
34:32	texel[3][1] index
47:45	texel[0][2] index
44:42	texel[1][2] index
41:40,55	texel[2][2] index
54:52	texel[3][2] index
51:49	texel[0][3] index
48,63:62	texel[1][3] index
61:59	texel[2][3] index
58:56	texel[3][3] index

The "ti" (table index) value from the compression block is used to select one of the columns in the table below.



Intensity modifier (im) table

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

The eight possible color values R_i are then computed from the 8 values in the column labeled im_i , where i ranges from 0 to 7:

For EAC_R11:

$$\text{if } (m == 0) R_i = R0 * 8 + 4 + im_i \text{ else } R_i = R0 * 8 + 4 + (im_i * m * 8)$$

Each value is clamped to the range [0,2047].

For EAC_SIGNED_R11:

$$\text{if } (m == 0) R_i = R0 * 8 + im_i \text{ else } R_i = R0 * 8 + (im_i * m * 8)$$

Each value is clamped to the range [-1023,1023].

Note that in the signed case, the $R0$ value is a signed, 2's complement value in the range [-127, 127]. Before being used in the above equations, an $R0$ value of -128 must be clamped to -127.

Finally, each texel red value is selected from the 8 possible values R_i using the 3-bit index for that texel. The green, blue, and alpha values are set to their default values.

The final value represents an 11-bit UNORM or SNORM as an unsigned/signed integer.

ETC2_RGB8_PTA and ETC2_SRGB8_PTA

The ETC2_RGB8_PTA format is similar to ETC2_RGB8 but eliminates the "individual" mode in favor of allowing a punch-through alpha. The "diff" bit from ETC2_RGB8 is renamed to "opaque" in this format, and the mode selection behaves as if the "diff" bit is always 1, making the "individual" mode inaccessible for these formats.

An alpha value of either 0 or 255 (representing 0.0 or 1.0) is possible with this format. If alpha is determined to be zero, the three other channels are also forced to zero, regardless of what value the normal decompression algorithm would have produced.

Differential Mode

In differential mode, if the opaque bit is set, the luminance table for ETC2_RGB8 is used. If the opaque bit is not set, the following luminance table is used (note that rows 0 and 2 have been zeroed out, otherwise the table is the same):

Luminance Table for opaque bit not set

	0	1	2	3	4	5	6	7
0	0	0	0	0	0	0	0	0
1	8	17	29	42	60	80	106	183
2	0	0	0	0	0	0	0	0
3	-8	-17	-29	-42	-60	-80	-106	-183

For each texel, if the opaque bit is zero and the corresponding texel index is equal to 2, the alpha value is set to zero (and therefore RGB for that texel will also end up at zero). Otherwise alpha is set to 255 and RGB is the result of the normal decompression calculations.

T and H Modes

In both of these modes, if the opaque bit is zero and the texel index is equal to 2, the alpha value is set to zero (and therefore RGB will also end up at zero). Otherwise, alpha is set to 255.

Planar Mode

In planar mode, the opaque bit is ignored, and alpha is set to 255.

The ETC2_SRGB8_PTA format is decompressed as if it is ETC2_RGB8_PTA, then a conversion from the resulting RGB values to SRGB space is performed, with alpha remaining unchanged.

ETC2_EAC_RGBA8 and ETC2_EAC_SRGB8_A8

The ETC2_EAC_RGBA8 format is a combination of ETC2_RGB8 and EAC_R8. A 16-byte compression block represents each 4x4. The low-order 8 bytes are used to compute alpha (instead of red) using the EAC_R8 algorithm. The high-order 8 bytes are used to compute RGB using the ETC2_RGB8 algorithm. The EAC_R8 format differs from EAC_R11 as described below.

The ETC2_EAC_SRGB8_A8 format is decompressed as if it is ETC2_EAC_RGBA8, then a conversion from the resulting RGB values to SRGB space is performed, with alpha remaining unchanged.



EAC_R8 Format:

The EAC_R8 format used within these surface formats is identical to EAC_R11 described in an earlier section, except the procedure for computing the eight possible color values R_i is performed as follows:

$$R_i = R_0 + (i m_i * m)$$

Each value is clamped to the range [0,255].

EAC_RG11 and EAC_SIGNED_RG11

These formats compress UNORM/SNORM double-channel data using a 16-byte compression block representing a 4x4 block of texels. The texels are labeled as $\text{texel}[\text{row}][\text{column}]$ where both row and column range from 0 to 3. $\text{Texel}[0][0]$ is the upper left texel.

The 16-byte compression block is laid out as follows.

EAC_RG11 compression block layout

Bits	Description
63:56	G0[7:0]
55:52	Gm[3:0]
51:48	Gti[3:0]
47:45	texel[0][0] G index
44:42	texel[1][0] G index
41:39	texel[2][0] G index
38:36	texel[3][0] G index
35:33	texel[0][1] G index
32:30	texel[1][1] G index
29:27	texel[2][1] G index
26:24	texel[3][1] G index
23:21	texel[0][2] G index
20:18	texel[1][2] G index
17:15	texel[2][2] G index
14:12	texel[3][2] G index
11:9	texel[0][3] G index
8:6	texel[1][3] G index
5:3	texel[2][3] G index
66:64	texel[3][3] G index
63:56	R0[7:0]
55:52	Rm[3:0]
51:48	Rti[3:0]
47:45	texel[0][0] R index
44:42	texel[1][0] R index

Bits	Description
41:39	texel[2][0] R index
38:36	texel[3][0] R index
35:33	texel[0][1] R index
32:30	texel[1][1] R index
29:27	texel[2][1] R index
26:24	texel[3][1] R index
23:21	texel[0][2] R index
20:18	texel[1][2] R index
17:15	texel[2][2] R index
14:12	texel[3][2] R index
11:9	texel[0][3] R index
8:6	texel[1][3] R index
5:3	texel[2][3] R index
2:0	texel[3][3] R index

These compression formats are identical to the EAC_R11 and EAC_SIGNED_R11 formats, except that they supply two channels of output data, both red and green, from two independent 8-byte portions of the compression block. The low half of the compression block contains the red information, and the high half contains the green information. Blue and alpha channels are set to their default values.

Refer to the EAC_R11 and EAC_SIGNED_R11 specification for details on how the red and green channels are generated using the data in the compression block.

DXT/BC1-3 Texture Formats

Note that non-power-of-2 dimensioned maps may require the surface to be padded out to the next multiple of four texels - here the pad texels are not referenced by the device.

An 8-byte (QWord) block encoding can be used if the source texture contains no transparency (is opaque) or if the transparency can be specified by a one-bit alpha. A 16-byte (DQWord) block encoding can be used to support source textures that require more than one-bit alpha: here the 1st QWord is used to encode the texel alpha values, and the 2nd QWord is used to encode the texel color values.

These three types of formats are discussed in the following sections:

- Opaque and One-bit Alpha Textures (DXT1)
- Opaque Textures (DXT1_RGB)
- Textures with Alpha Channels (DXT2-5)

DXT2 and DXT3 are equivalent compression formats from the perspective of the hardware. The only difference between the two is the use of pre-multiplied alpha encoding, which does not affect hardware.

Likewise, DXT4 and DXT5 are the same compression formats with the only difference being the use of pre-multiplied alpha encoding.



Note that the surface formats DXT1-5 are referred to in the DirectX Specification as BC1-3. The mapping between formats is shown below:

- DXT1 => BC1
- DXT2/DXT3 => BC2
- DXT4/DXT5 => BC3

Programming Note	
Context:	DXT Texture Formats
<ul style="list-style-type: none">• Any single texture must specify that its data is stored as 64 or 128 bits per group of 16 texels. If 64-bit blocks--that is, format DXT1--are used for the texture, it is possible to mix the opaque and one-bit alpha formats on a per-block basis within the same texture. In other words, the comparison of the unsigned integer magnitude of color_0 and color_1 is performed uniquely for each block of 16 texels.• When 128-bit blocks are used, then the alpha channel must be specified in either explicit (format DXT2 or DXT3) or interpolated mode (format DXT4 or DXT5) for the entire texture. Note that as with color, once interpolated mode is selected then either 8 interpolated alphas or 6 interpolated alphas mode can be used on a block-by-block basis. Again the magnitude comparison of alpha_0 and alpha_1 is done uniquely on a block-by-block basis.	

Opaque and One-bit Alpha Textures (DXT1/BC1)

Texture format DXT1 is for textures that are opaque or have a single transparent color. For each opaque or one-bit alpha block, two 16-bit R5G6B5 values and a 4x4 bitmap with 2-bits-per-pixel are stored. This totals 64 bits (1 QWord) for 16 texels, or 4-bits-per-texel.

In the block bitmap, there are two bits per texel to select between the four colors, two of which are stored in the encoded data. The other two colors are derived from these stored colors by linear interpolation.

The one-bit alpha format is distinguished from the opaque format by comparing the two 16-bit color values stored in the block. They are treated as unsigned integers. If the first color is greater than the second, it implies that only opaque texels are defined. This means four colors will be used to represent the texels. In four-color encoding, there are two derived colors and all four colors are equally distributed in RGB color space. This format is analogous to R5G6B5 format. Otherwise, for one-bit alpha transparency, three colors are used and the fourth is reserved to represent transparent texels. Note that the color blocks in DXT2-5 formats strictly use four colors, as the alpha values are obtained from the alpha block

In three-color encoding, there is one derived color and the fourth two-bit code is reserved to indicate a transparent texel (alpha information). This format is analogous to A1R5G5B5, where the final bit is used for encoding the alpha mask.

The following piece of pseudo-code illustrates the algorithm for deciding whether three- or four-color encoding is selected:

```

if (color_0 > color_1)
{
  // Four-color block: derive the other two colors.
  // 00 = color_0, 01 = color_1, 10 = color_2, 11 = color_3
  // These two bit codes correspond to the 2-bit fields
  // stored in the 64-bit block.
  color_2 = (2 * color_0 + color_1) / 3;
  color_3 = (color_0 + 2 * color_1) / 3;
}
else
{
  // Three-color block: derive the other color.
  // 00 = color_0, 01 = color_1, 10 = color_2,
  // 11 = transparent.
  // These two bit codes correspond to the 2-bit fields
  // stored in the 64-bit block.
  color_2 = (color_0 + color_1) / 2;
  color_3 = transparent;
}

```

The following tables show the memory layout for the 8-byte block. It is assumed that the first index corresponds to the y-coordinate and the second corresponds to the x-coordinate. For example, Texel[1][2] refers to the texture map pixel at (x,y) = (2,1).

Here is the memory layout for the 8-byte (64-bit) block:

Word Address	16-bit Word
0	Color_0
1	Color_1
2	Bitmap Word_0
3	Bitmap Word_1

Color_0 and Color_1 (colors at the two extremes) are laid out as follows:

Bits	Color
15:11	Red color component
10:5	Green color component
4:0	Blue color component

Bits	Texel
1:0 (LSB)	Texel[0][0]
3:2	Texel[0][1]
5:4	Texel[0][2]
7:6	Texel[0][3]
9:8	Texel[1][0]
11:10	Texel[1][1]
13:12	Texel[1][2]



Bits	Texel
15:14	Texel[1][3]

Bitmap Word_1 is laid out as follows:

Bits	Texel
1:0 (LSB)	Texel[2][0]
3:2	Texel[2][1]
5:4	Texel[2][2]
7:6	Texel[2][3]
9:8	Texel[3][0]
11:10	Texel[3][1]
13:12	Texel[3][2]
15:14 (MSB)	Texel[3][3]

Example of Opaque Color Encoding

As an example of opaque encoding, we will assume that the colors red and black are at the extremes. We will call red `color_0` and black `color_1`. There will be four interpolated colors that form the uniformly distributed gradient between them. To determine the values for the 4x4 bitmap, the following calculations are used:

```
00 ? color_0
01 ? color_1
10 ? 2/3 color_0 + 1/3 color_1
11 ? 1/3 color_0 + 2/3 color_1
```

Example of One-bit Alpha Encoding

This format is selected when the unsigned 16-bit integer, `color_0`, is less than the unsigned 16-bit integer, `color_1`. An example of where this format could be used is leaves on a tree to be shown against a blue sky. Some texels could be marked as transparent while three shades of green are still available for the leaves. Two of these colors fix the extremes, and the third color is an interpolated color.

The bitmap encoding for the colors and the transparency is determined using the following calculations:

```
00 ? color_0
01 ? color_1
10 ? 1/2 color_0 + 1/2 color_1
11 ? Transparent
```

Opaque Textures (DXT1_RGB)

Texture format DXT1_RGB is identical to DXT1, with the exception that the One-bit Alpha encoding is removed. Color 0 and Color 1 are not compared, and the resulting texel color is derived strictly from the Opaque Color Encoding. The alpha channel defaults to 1.0.

Programming Note	
Context:	Opaque Textures (DXT1_RGB)
The behavior of this format is not compliant with the OGL spec.	

Compressed Textures with Alpha Channels (DXT2-5 / BC2-3)

There are two ways to encode texture maps that exhibit more complex transparency. In each case, a block that describes the transparency precedes the 64-bit block already described for DXT1. The transparency is either represented as a 4x4 bitmap with four bits per pixel (explicit encoding), or with fewer bits and linear interpolation analogous to what is used for color encoding.

The transparency block and the color block are laid out as follows:

Word Address	64-bit Block
3:0	Transparency block
7:4	Previously described 64-bit block

Explicit Texture Encoding

For explicit texture encoding (DXT2 and DXT3 formats), the alpha components of the texels that describe transparency are encoded in a 4x4 bitmap with 4 bits per texel. These 4 bits can be achieved through a variety of means such as dithering or by simply using the 4 most significant bits of the alpha data. However, they are produced, they are used just as they are, without any form of interpolation.

Note: DirectDraw's compression method uses the 4 most significant bits.

The following tables illustrate how the alpha information is laid out in memory, for each 16-bit word.

This is the layout for Word 0:

Bits	Alpha
3:0 (LSB)	[0][0]
7:4	[0][1]
11:8	[0][2]
15:12 (MSB)	[0][3]



This is the layout for Word 1:

Bits	Alpha
3:0 (LSB)	[1][0]
7:4	[1][1]
11:8	[1][2]
15:12 (MSB)	[1][3]

This is the layout for Word 2:

Bits	Alpha
3:0 (LSB)	[2][0]
7:4	[2][1]
11:8	[2][2]
15:12 (MSB)	[2][3]

This is the layout for Word 3:

Bits	Alpha
3:0 (LSB)	[3][0]
7:4	[3][1]
11:8	[3][2]
15:12 (MSB)	[3][3]

Three-Bit Linear Alpha Interpolation

The encoding of transparency for the DXT4 and DXT5 formats is based on a concept similar to the linear encoding used for color. Two 8-bit alpha values and a 4x4 bitmap with three bits per pixel are stored in the first eight bytes of the block. The representative alpha values are used to interpolate intermediate alpha values. Additional information is available in the way the two alpha values are stored. If alpha_0 is greater than alpha_1, then six intermediate alpha values are created by the interpolation. Otherwise, four intermediate alpha values are interpolated between the specified alpha extremes. The two additional implicit alpha values are 0 (fully transparent) and 255 (fully opaque).

The following pseudo-code illustrates this algorithm:

```
// 8-alpha or 6-alpha block?
if (alpha_0 > alpha_1) {
    // 8-alpha block: derive the other 6 alphas.
    // 000 = alpha_0, 001 = alpha_1, others are interpolated
    alpha_2 = (6 * alpha_0 + alpha_1) / 7; // Bit code 010
    alpha_3 = (5 * alpha_0 + 2 * alpha_1) / 7; // Bit code 011
    alpha_4 = (4 * alpha_0 + 3 * alpha_1) / 7; // Bit code 100
    alpha_5 = (3 * alpha_0 + 4 * alpha_1) / 7; // Bit code 101
    alpha_6 = (2 * alpha_0 + 5 * alpha_1) / 7; // Bit code 110
    alpha_7 = (alpha_0 + 6 * alpha_1) / 7; // Bit code 111
}
else {
    // 6-alpha block: derive the other alphas.
    // 000 = alpha_0, 001 = alpha_1, others are interpolated
    alpha_2 = (4 * alpha_0 + alpha_1) / 5; // Bit code 010
    alpha_3 = (3 * alpha_0 + 2 * alpha_1) / 5; // Bit code 011
}
```

```

alpha_4 = (2 * alpha_0 + 3 * alpha_1) / 5; // Bit code 100
alpha_5 = (alpha_0 + 4 * alpha_1) / 5;    // Bit code 101
alpha_6 = 0;                               // Bit code 110
alpha_7 = 255;                             // Bit code 111
}

```

The memory layout of the alpha block is as follows:

Byte	Alpha
0	Alpha_0
1	Alpha_1
2	[0][2] (2 LSBs), [0][1], [0][0]
3	[1][1] (1 LSB), [1][0], [0][3], [0][2] (1 MSB)
4	[1][3], [1][2], [1][1] (2 MSBs)
5	[2][2] (2 LSBs), [2][1], [2][0]
6	[3][1] (1 LSB), [3][0], [2][3], [2][2] (1 MSB)
7	[3][3], [3][2], [3][1] (2 MSBs)

BC4

These formats (BC4_UNORM and BC4_SNORM) compresses single-component UNORM or SNORM data. An 8-byte compression block represents a 4x4 block of texels. The texels are labeled as texel[row][column] where both row and column range from 0 to 3. Texel[0][0] is the upper left texel.

The 8-byte compression block is laid out as follows:

Bit	Description
7:0	red_0
15:8	red_1
18:16	texel[0][0] bit code
21:19	texel[0][1] bit code
24:22	texel[0][2] bit code
27:25	texel[0][3] bit code
30:28	texel[1][0] bit code
33:31	texel[1][1] bit code
36:34	texel[1][2] bit code
39:37	texel[1][3] bit code
42:40	texel[2][0] bit code
45:43	texel[2][1] bit code
48:46	texel[2][2] bit code
51:49	texel[2][3] bit code
54:52	texel[3][0] bit code
57:55	texel[3][1] bit code



Bit	Description
60:58	texel[3][2] bit code
63:61	texel[3][3] bit code

There are two interpolation modes, chosen based on which reference color is larger. The first mode has the two reference colors plus six equal-spaced interpolated colors between the reference colors, chosen based on the three-bit code for that texel. The second mode has the two reference colors plus four interpolated colors, chosen by six of the three-bit codes. The remaining two codes select min and max values for the colors. The values of red_0 through red_7 are computed as follows:

```

red_0 = red_0; // bit code 000
red_1 = red_1; // bit code 001
if (red_0 > red_1) {
    red_2 = (6 * red_0 + 1 * red_1) / 7; // bit code 010
    red_3 = (5 * red_0 + 2 * red_1) / 7; // bit code 011
    red_4 = (4 * red_0 + 3 * red_1) / 7; // bit code 100
    red_5 = (3 * red_0 + 4 * red_1) / 7; // bit code 101
    red_6 = (2 * red_0 + 5 * red_1) / 7; // bit code 110
    red_7 = (1 * red_0 + 6 * red_1) / 7; // bit code 111
}
else {
    red_2 = (4 * red_0 + 1 * red_1) / 5; // bit code 010
    red_3 = (3 * red_0 + 2 * red_1) / 5; // bit code 011
    red_4 = (2 * red_0 + 3 * red_1) / 5; // bit code 100
    red_5 = (1 * red_0 + 4 * red_1) / 5; // bit code 101
    red_6 = UNORM ? 0.0 : -1.0; // bit code 110 (0 for UNORM, -1 for SNORM)
    red_7 = 1.0; // bit code 111
}

```

BC5

These formats (BC5_UNORM and BC5_SNORM) compresses dual-component UNORM or SNORM data. A 16-byte compression block represents a 4x4 block of texels. The texels are labeled as texel[row][column] where both row and column range from 0 to 3. Texel[0][0] is the upper left texel.

The 16-byte compression block is laid out as follows:

Bit	Description
7:0	red_0
15:8	red_1
18:16	texel[0][0] red bit code
21:19	texel[0][1] red bit code
24:22	texel[0][2] red bit code
27:25	texel[0][3] red bit code
30:28	texel[1][0] red bit code
33:31	texel[1][1] red bit code
36:34	texel[1][2] red bit code
39:37	texel[1][3] red bit code
42:40	texel[2][0] red bit code

Bit	Description
45:43	texel[2][1] red bit code
48:46	texel[2][2] red bit code
51:49	texel[2][3] red bit code
54:52	texel[3][0] red bit code
57:55	texel[3][1] red bit code
60:58	texel[3][2] red bit code
63:61	texel[3][3] red bit code
71:64	green_0
79:72	green_1
82:80	texel[0][0] green bit code
85:83	texel[0][1] green bit code
88:86	texel[0][2] green bit code
91:89	texel[0][3] green bit code
94:92	texel[1][0] green bit code
97:95	texel[1][1] green bit code
100:98	texel[1][2] green bit code
103:101	texel[1][3] green bit code
106:104	texel[2][0] green bit code
109:107	texel[2][1] green bit code
112:110	texel[2][2] green bit code
115:113	texel[2][3] green bit code
118:116	texel[3][0] green bit code
121:119	texel[3][1] green bit code
124:122	texel[3][2] green bit code
127:125	texel[3][3] green bit code

There are two interpolation modes, chosen based on which reference color is larger. The first mode has the two reference colors plus six equal-spaced interpolated colors between the reference colors, chosen based on the three-bit code for that texel. The second mode has the two reference colors plus four interpolated colors, chosen by six of the three-bit codes. The remaining two codes select min and max values for the colors. The values of red_0 through red_7 are computed as follows:



```
red_0 = red_0; // bit code 000
red_1 = red_1; // bit code 001
if (red_0 > red_1) {
    red_2 = (6 * red_0 + 1 * red_1) / 7; // bit code 010
    red_3 = (5 * red_0 + 2 * red_1) / 7; // bit code 011
    red_4 = (4 * red_0 + 3 * red_1) / 7; // bit code 100
    red_5 = (3 * red_0 + 4 * red_1) / 7; // bit code 101
    red_6 = (2 * red_0 + 5 * red_1) / 7; // bit code 110
    red_7 = (1 * red_0 + 6 * red_1) / 7; // bit code 111
}
else {
    red_2 = (4 * red_0 + 1 * red_1) / 5; // bit code 010
    red_3 = (3 * red_0 + 2 * red_1) / 5; // bit code 011
    red_4 = (2 * red_0 + 3 * red_1) / 5; // bit code 100
    red_5 = (1 * red_0 + 4 * red_1) / 5; // bit code 101
    red_6 = UNORM ? 0.0 : -1.0; // bit code 110 (0 for UNORM, -1 for SNORM)
    red_7 = 1.0; // bit code 111
}
```

The same calculations are done for green, using the corresponding reference colors and bit codes.

BC6H

These formats (BC6H_UF16 and BC6H_SF16) compresses 3-channel images with high dynamic range (> 8 bits per channel). BC6H supports floating point denorms but there is no support for INF and NaN, other than with BC6H_SF16 -INF is supported. The alpha channel is not included, thus alpha is returned at its default value.

The BC6H block is 16 bytes and represents a 4x4 block of texels. The texels are labeled as texel[row][column] where both row and column range from 0 to 3. Texel[0][0] is the upper left texel. BC6H has 14 different modes, the mode that the block is in is contained in the least significant bits (either 2 or 5 bits).

The basic scheme consists of interpolating colors along either one or two lines, with per-texel indices indicating which color along the line is chosen for each texel. If a two-line mode is selected, one of 32 partition sets is indicated which selects which of the two lines each texel is assigned to.

Field Definition

There are 14 possible modes for a BC6H block, the format of each is indicated in the 14 tables below. The mode is selected by the unique mode bits specified in each table. The first 10 modes use two lines ("TWO"), and the last 4 use one line ("ONE"). The difference between the various two-line and one-line modes is with the precision of the first endpoint and the number of bits used to store delta values for the remaining endpoints. Two modes (9 and 10) specify each endpoint as an original value rather than using the deltas (these are indicated as having no delta values).

The endpoints values and deltas are indicated in the tables using a two-letter name. The first letter is "r", "g", or "b" indicating the color channel. The second letter is "w", "x", "y", or "z" indicating which of the four endpoints. The first line has endpoints "w" and "x", with "w" being the endpoint that is fully specified (i.e. not as a delta). The second line has endpoints "y" and "z". Modes using ONE mode do not have endpoints "y" and "z" as they have only one line.

In addition to the mode and endpoint data, TWO blocks contain a 5-bit "partition" which selects one of the partition sets, and a 46-bit set of indices. ONE blocks contain a 63-bit set of indices. These are described in more detail below.

Mode 0: (TWO) Red, Green, Blue: 10-bit endpoint, 5-bit deltas

Bit	Description
1:0	mode = 00
2	gy[4]
3	by[4]
4	bz[4]
14:5	rw[9:0]
24:15	gw[9:0]
34:25	bw[9:0]
39:35	rx[4:0]
40	gz[4]
44:41	gy[3:0]
49:45	gx[4:0]
50	bz[0]
54:51	gz[3:0]
59:55	bx[4:0]
60	bz[1]
64:61	by[3:0]
69:65	ry[4:0]
70	bz[2]
75:71	rz[4:0]
76	bz[3]
81:77	partition
127:82	indices

Mode 1: (TWO) Red, Green, Blue: 7-bit endpoint, 6-bit deltas

Bit	Description
1:0	mode = 01
2	gy[5]
3	gz[4]
4	gz[5]
11:5	rw[6:0]
12	bz[0]
13	bz[1]
14	by[4]

Bit	Description
21:15	gw[6:0]
22	by[5]
23	bz[2]
24	gy[4]
31:25	bw[6:0]
32	bz[3]
33	bz[5]
34	bz[4]
40:35	rx[5:0]
44:41	gy[3:0]
50:45	gx[5:0]
54:51	gz[3:0]
60:55	bx[5:0]
64:61	by[3:0]
70:65	ry[5:0]
76:71	rz[5:0]
81:77	partition
127:82	indices

Mode 2: (TWO) Red: 11-bit endpoint, 5-bit deltas

Green, Blue: 11-bit endpoint, 4-bit deltas

Bit	Description
4:0	mode = 00010
14:5	rw[9:0]
24:15	gw[9:0]
34:25	bw[9:0]
39:35	rx[4:0]
40	rw[10]
44:41	gy[3:0]
48:45	gx[3:0]
49	gw[10]
50	bz[0]
54:51	gz[3:0]
58:55	bx[3:0]
59	bw[10]
60	bz[1]
64:61	by[3:0]

Bit	Description
69:65	ry[4:0]
70	bz[2]
75:71	rz[4:0]
76	bz[3]
81:77	partition
127:82	indices

Mode 3: (TWO) Red, Blue: 11-bit endpoint, 4-bit deltas

Green: 11-bit endpoint, 5-bit deltas

Bit	Description
4:0	mode = 00110
14:5	rw[9:0]
24:15	gw[9:0]
34:25	bw[9:0]
38:35	rx[3:0]
39	rw[10]
40	gz[4]
44:41	gy[3:0]
49:45	gx[4:0]
50	gw[10]
54:51	gz[3:0]
58:55	bx[3:0]
59	bw[10]
60	bz[1]
64:61	by[3:0]
68:65	ry[3:0]
69	bz[0]
70	bz[2]
74:71	rz[3:0]
75	gy[4]
76	bz[3]
81:77	partition
127:82	indices



Mode 4: (TWO) Red, Green: 11-bit endpoint, 4-bit deltas

Blue: 11-bit endpoint, 5-bit deltas

Bit	Description
4:0	mode = 01010
14:5	rw[9:0]
24:15	gw[9:0]
34:25	bw[9:0]
38:35	rx[3:0]
39	rw[10]
40	by[4]
44:41	gy[3:0]
48:45	gx[3:0]
49	gw[10]
50	bz[0]
54:51	gz[3:0]
59:55	bx[4:0]
60	bw[10]
64:61	by[3:0]
68:65	ry[3:0]
69	bz[1]
70	bz[2]
74:71	rz[3:0]
75	bz[4]
76	bz[3]
81:77	partition
127:82	indices

Mode 5: (TWO) Red, Green, Blue: 9-bit endpoint, 5-bit deltas

Bit	Description
4:0	mode = 01110
13:5	rw[8:0]
14	by[4]
23:15	gw[8:0]
24	gy[4]
33:25	bw[8:0]
34	bz[4]
39:35	rx[4:0]
40	gz[4]

Bit	Description
44:41	gy[3:0]
49:45	gx[4:0]
50	bz[0]
54:51	gz[3:0]
59:55	bx[4:0]
60	bz[1]
64:61	by[3:0]
69:65	ry[4:0]
70	bz[2]
75:71	rz[4:0]
76	bz[3]
81:77	partition
127:82	indices

Mode 6: (TWO) Red: 8-bit endpoint, 6-bit deltas

Green, Blue: 8-bit endpoint, 5-bit deltas

Bit	Description
4:0	mode = 10010
12:5	rw[7:0]
13	gz[4]
14	by[4]
22:15	gw[7:0]
23	bz[2]
24	gy[4]
32:25	bw[7:0]
33	bz[3]
34	bz[4]
40:35	rx[5:0]
44:41	gy[3:0]
49:45	gx[4:0]
50	bz[0]
54:51	gz[3:0]
59:55	bx[4:0]
60	gz[1]
64:61	by[3:0]
70:65	ry[5:0]
76:71	rz[5:0]



Bit	Description
81:77	partition
127:82	indices

Mode 7: (TWO) Red, Blue: 8-bit endpoint, 5-bit deltas

Green: 8-bit endpoint, 6-bit deltas

Bit	Description
4:0	mode = 10110
12:5	rw[7:0]
13	bz[0]
14	by[4]
22:15	gw[7:0]
23	gy[5]
24	gy[4]
32:25	bw[7:0]
33	gz[5]
34	bz[4]
39:35	rx[4:0]
40	gz[4]
44:41	gy[3:0]
50:45	gx[5:0]
54:51	gz[3:0]
59:55	bx[4:0]
60	bz[1]
64:61	by[3:0]
69:65	ry[4:0]
70	bz[2]
75:71	rz[4:0]
76	bz[3]
81:77	partition
127:82	indices

Mode 8: (TWO) Red, Green: 8-bit endpoint, 5-bit deltas

Blue: 8-bit endpoint, 6-bit deltas

Bit	Description
4:0	mode = 11010
12:5	rw[7:0]
13	bz[1]

Bit	Description
14	by[4]
22:15	gw[7:0]
23	by[5]
24	gy[4]
32:25	bw[7:0]
33	bz[5]
34	bz[4]
39:35	rx[4:0]
40	gz[4]
44:41	gy[3:0]
49:45	gx[4:0]
50	bz[0]
54:51	gz[3:0]
60:55	bx[5:0]
64:61	by[3:0]
69:65	ry[4:0]
70	bz[2]
75:71	rz[4:0]
76	bz[3]
81:77	partition
127:82	indices

Mode 9: (TWO) Red, Green, Blue: 6-bit endpoints for all four, no deltas

Bit	Description
4:0	mode = 11110
10:5	rw[5:0]
11	gz[4]
12	bz[0]
13	bz[1]
14	by[4]
20:15	gw[5:0]
21	gy[5]
22	by[5]
23	bz[2]
24	gy[4]
30:25	bw[5:0]
31	gz[5]



Bit	Description
32	bz[3]
33	bz[5]
34	bz[4]
40:35	rx[5:0]
44:41	gy[3:0]
50:45	gx[5:0]
54:51	gz[3:0]
60:55	bx[5:0]
64:61	by[3:0]
70:65	ry[5:0]
76:71	rz[5:0]
81:77	partition
127:82	indices

Mode 10: (ONE) Red, Green, Blue: 10-bit endpoints for both, no deltas

Bit	Description
4:0	mode = 00011
14:5	rw[9:0]
24:15	gw[9:0]
34:25	bw[9:0]
44:35	rx[9:0]
54:45	gx[9:0]
64:55	bx[9:0]
127:65	indices

Mode 11: (ONE) Red, Green, Blue: 11-bit endpoints, 9-bit deltas

Bit	Description
4:0	mode = 00111
14:5	rw[9:0]
24:15	gw[9:0]
34:25	bw[9:0]
43:35	rx[8:0]
44	rw[10]
53:45	gx[8:0]
54	gw[10]
63:55	bx[8:0]
64	bw[10]
127:65	indices

Mode 12: (ONE) Red, Green, Blue: 12-bit endpoints, 8-bit deltas

Bit	Description
4:0	mode = 01011
14:5	rw[9:0]
24:15	gw[9:0]
34:25	bw[9:0]
42:35	rx[7:0]
43	rw[11]
44	rw[10]
52:45	gx[7:0]
53	gw[11]
54	gw[10]
62:55	bx[7:0]
63	bw[11]
64	bw[10]
127:65	indices

Mode 13: (ONE) Red, Green, Blue: 16-bit endpoints, 4-bit deltas

Bit	Description
4:0	mode = 01111
14:5	rw[9:0]
24:15	gw[9:0]
34:25	bw[9:0]
38:35	rx[3:0]
39	rw[15]
40	rw[14]
41	rw[13]
42	rw[12]
43	rw[11]
44	rw[10]
48:45	gx[3:0]
49	gw[15]
50	gw[14]
51	gw[13]
52	gw[12]
53	gw[11]
54	gw[10]
58:55	bx[3:0]



Bit	Description
59	bw[15]
60	bw[14]
61	bw[13]
62	bw[12]
63	bw[11]
64	bw[10]
127:65	indices

Undefined mode values (10011, 10111, 11011, and 11111) return zero in the RGB channels.

The "indices" fields are defined as follows:

TWO mode *indices* field with fix-up index [1] at texel[3][3]

Bit	Description
83:82	texel[0][0] index
86:84	texel[0][1] index
89:87	texel[0][2] index
92:90	texel[0][3] index
95:93	texel[1][0] index
98:96	texel[1][1] index
101:99	texel[1][2] index
104:102	texel[1][3] index
107:105	texel[2][0] index
110:108	texel[2][1] index
113:111	texel[2][2] index
116:114	texel[2][3] index
119:117	texel[3][0] index
122:120	texel[3][1] index
125:123	texel[3][2] index
127:126	texel[3][3] index

TWO mode *indices* field with fix-up index [1] at texel[0][2]

Bit	Description
83:82	texel[0][0] index
86:84	texel[0][1] index
88:87	texel[0][2] index
91:89	texel[0][3] index
94:92	texel[1][0] index
97:95	texel[1][1] index

Bit	Description
100:98	texel[1][2] index
103:101	texel[1][3] index
106:104	texel[2][0] index
109:107	texel[2][1] index
112:110	texel[2][2] index
115:113	texel[2][3] index
118:116	texel[3][0] index
121:119	texel[3][1] index
124:122	texel[3][2] index
127:125	texel[3][3] index

TWO mode *indices* field with fix-up index [1] at texel[2][0]

Bit	Description
83:82	texel[0][0] index
86:84	texel[0][1] index
89:87	texel[0][2] index
92:90	texel[0][3] index
95:93	texel[1][0] index
98:96	texel[1][1] index
101:99	texel[1][2] index
104:102	texel[1][3] index
106:105	texel[2][0] index
109:107	texel[2][1] index
112:110	texel[2][2] index
115:113	texel[2][3] index
118:116	texel[3][0] index
121:119	texel[3][1] index
124:122	texel[3][2] index
127:125	texel[3][3] index

ONE mode *indices* field

Bit	Description
67:65	texel[0][0] index
71:68	texel[0][1] index
75:72	texel[0][2] index
79:76	texel[0][3] index
83:80	texel[1][0] index
87:84	texel[1][1] index



Bit	Description
91:88	texel[1][2] index
95:92	texel[1][3] index
99:96	texel[2][0] index
103:100	texel[2][1] index
107:104	texel[2][2] index
111:108	texel[2][3] index
115:112	texel[3][0] index
119:116	texel[3][1] index
123:120	texel[3][2] index
127:124	texel[3][3] index

Endpoint Computation

The endpoints can be defined in many different ways, as shown above. This section describes how the endpoints are computed from the bits in the compression block. The method used depends on whether the BC6H format is signed (BC6H_SF16) or unsigned (BC6H_UF16).

First, each channel (RGB) of each endpoint is extended to 16 bits. Each is handled identically and independently, however in some modes different channels have different incoming precision which must be accounted for. The following rules are employed:

- If the format is BC6H_SF16 or the endpoint is a delta value, the value is sign-extended to 16 bits
- For all other cases, the value is zero-extended to 16 bits

If there are no endpoints that are delta values, endpoint computation is complete. For endpoints that are delta values, the next step involves computing the absolute endpoint. The "w" endpoint is always absolute and acts as a base value for the other three endpoints. Each channel is handled identically and independently.

$$\begin{aligned}x &= w + x \\y &= w + y \\z &= w + z\end{aligned}$$

The above is performed using 16-bit integer arithmetic. Overflows beyond 16 bits are ignored (any resulting high bits are dropped).

Palette Color Computation

The next step involves computing the color palette values that provide the available values for each texel's color. The color palette for each line consists of the two endpoint colors plus 6 (TWO mode) or 14 (ONE mode) interpolated colors. Again, each channel is processed independently.

First the endpoints are unquantized, with each channel of each endpoint being processed independently. The number of bits in the original base w value represents the precision of the endpoints. The input

endpoint is called e , and the resulting endpoints are represented as 17-bit signed integers and called e' below.

For the BC6H_UF16 format:

- if the precision is already 16 bits, $e' = e$
- if $e = 0$, $e' = 0$
- if e is the maximum representable in the precision, $e' = 0xFFFF$
- otherwise, $e' = ((e \ll 16) + 0x8000) \gg \text{precision}$

For the BC6H_SF16 format, the value is treated as sign magnitude. The sign is not changed, e' and e refer only to the magnitude portion:

- if the precision is already 16 bits, $e' = e$
- if $e = 0$, $e' = 0$
- if e is the maximum representable in the precision, $e' = 0x7FFF$
- otherwise, $e' = ((e \ll 15) + 0x4000) \gg (\text{precision} - 1)$

Next, the palette values are generated using predefined weights, using the tables below:

$$\text{palette}[i] = (w' * (64 - \text{weight}[i]) + x' * \text{weight}[i] + 32) \gg 6$$

TWO mode weights:

palette index	0	1	2	3	4	5	6	7
weight	0	9	18	27	37	46	55	64

ONE mode weights:

palette index	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
weight	0	4	9	13	17	21	26	30	34	38	43	47	51	55	60	64

The two end palette indices are equal to the two endpoints given that the weights are 0 and 64. In the above equation w' and x' represent the endpoints e' computed in the previous step corresponding to w and x , respectively. For the second line in TWO mode, w and x are replaced with y and z .

The final step in computing the palette colors is to rescale the final results. For BC6H_UF16 format, the values are multiplied by 31/64. For BC6H_SF16, the values are multiplied by 31/32, treating them as sign magnitude. These final 16-bit results are ultimately treated as 16-bit floats.

Texel Selection

The final step is to select the appropriate palette index for each texel. This index then selects the 16-bit per channel palette value, which is re-interpreted as a 16-bit floating point result for input into the filter. This procedure differs depending on whether the mode is TWO or ONE.



ONE Mode

In ONE mode, there is only one set of palette colors, but the "indices" field is 63 bits. This field consists of a 4-bit palette index for each of the 16 texels, with the exception of the texel at [0][0] which has only 3 bits, the missing high bit being set to zero.

TWO Mode

32 partitions are defined for TWO, which are defined below. Each of the 32 cases shows the 4x4 block of texels, and is indexed by adding its hexadecimal row number (00-1C) to its column number (0-3). Each texel in the 4x4 block indicates whether that texel is part of line 0 (endpoints w and x) or line 1 (endpoints y and z). Each case has one texel each of "[0]" and "[1]", the index that this is at is termed the "fix-up index". These texels have one less bit in the index.

	0				1				2				3			
00	[0]	0	1	1	[0]	0	0	1	[0]	1	1	1	[0]	0	0	1
	0	0	1	1	0	0	0	1	0	1	1	1	0	0	1	1
	0	0	1	1	0	0	0	1	0	1	1	1	0	0	1	1
	0	0	1	[1]	0	0	0	[1]	0	1	1	[1]	0	1	1	[1]
04	[0]	0	0	0	[0]	0	1	1	[0]	0	0	1	[0]	0	0	0
	0	0	0	1	0	1	1	1	0	0	1	1	0	0	0	1
	0	0	0	1	0	1	1	1	0	1	1	1	0	0	1	1
	0	0	1	[1]	1	1	1	[1]	1	1	1	[1]	0	1	1	[1]
08	[0]	0	0	0	[0]	0	1	1	[0]	0	0	0	[0]	0	0	0
	0	0	0	0	0	1	1	1	0	0	0	1	0	0	0	0
	0	0	0	1	1	1	1	1	0	1	1	1	0	0	0	1
	0	0	1	[1]	1	1	1	[1]	1	1	1	[1]	0	1	1	[1]
0C	[0]	0	0	1	[0]	0	0	0	[0]	0	0	0	[0]	0	0	0
	0	1	1	1	0	0	0	0	1	1	1	1	0	0	0	0
	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0
	1	1	1	[1]	1	1	1	[1]	1	1	1	[1]	1	1	1	[1]
10	[0]	0	0	0	[0]	1	[1]	1	[0]	0	0	0	[0]	1	[1]	1
	1	0	0	0	0	0	0	1	0	0	0	0	0	0	1	1
	1	1	1	0	0	0	0	0	[1]	0	0	0	0	0	0	1
	1	1	1	[1]	0	0	0	0	1	1	1	0	0	0	0	0
14	[0]	0	[1]	1	[0]	0	0	0	[0]	0	0	0	[0]	1	1	1
	0	0	0	1	1	0	0	0	0	0	0	0	0	0	1	1
	0	0	0	0	[1]	1	0	0	[1]	0	0	0	0	0	1	1
	0	0	0	0	1	1	1	0	1	1	0	0	0	0	0	[1]
18	[0]	0	[1]	1	[0]	0	0	0	[0]	1	[1]	0	[0]	0	[1]	1
	0	0	0	1	1	0	0	0	0	1	1	0	0	1	1	0
	0	0	0	1	[1]	0	0	0	0	1	1	0	0	1	1	0

	0				1				2				3			
	0	0	0	0	1	1	0	0	0	1	1	0	1	1	0	0
1C	[0]	0	0	1	[0]	0	0	0	[0]	1	[1]	1	[0]	0	[1]	1
	0	1	1	1	1	1	1	1	0	0	0	1	1	0	0	1
	[1]	1	1	0	[1]	1	1	1	1	0	0	0	1	0	0	1
	1	0	0	0	0	0	0	0	1	1	1	0	1	1	0	0

The 46-bit "indices" field consists of a 3-bit palette index for each of the 16 texels, with the exception of the bracketed texels that have only two bits each. The high bit of these texels is set to zero.

BC7

These formats (BC7_UNORM and BC7_UNORM_SRGB) compresses 3-channel and 4-channel fixed point images.

The BC7 block is 16 bytes and represents a 4x4 block of texels. The texels are labeled as texel[row][column] where both row and column range from 0 to 3. Texel[0][0] is the upper left texel. BC7 has 8 different modes, the mode that the block is in is contained in the least significant bits (1-8 bits depending on mode).

The basic scheme consists of interpolating colors and alpha in some modes along either one, two, or three lines, with per-texel indices indicating which color/alpha along the line is chosen for each texel. If a two- or three-line mode is selected, one of 64 partition sets is indicated which selects which of the two lines each texel is assigned to, although some modes are limited to the first 16 partition sets. In the color-only modes, alpha is always returned at its default value of 1.0.

Some modes contain the following fields:

- **P-bits.** These represent shared LSB for all components of the endpoint, which increases the endpoint precision by one bit. In some cases both endpoints of a line share a P-bit.
- **Rotation bits.** For blocks with separate color and alpha, this 2-bit field allows selection of which of the four components has its own indexes (scalar) vs. the other three components (vector).
- **Index selector.** This 1-bit field selects whether the scalar or vector components uses the 3-bit index vs. the 2-bit index.

Field Definition

There are 8 possible modes for a BC7 block, the format of each is indicated in the 8 tables below. The mode is selected by the unique mode bits specified in each table. Each mode has particular characteristics described at the top of the table.

Mode 0: Color only, 3 lines (THREE), 4-bit endpoints with one P-bit per endpoint, 3-bit indices, 16 partitions

Bit	Description
0	mode = 0



Bit	Description
4:1	partition
8:5	R0
12:9	R1
16:13	R2
20:17	R3
24:21	R4
28:25	R5
32:29	G0
36:33	G1
40:37	G2
44:41	G3
48:45	G4
52:49	G5
56:53	B0
60:57	B1
64:61	B2
68:65	B3
72:69	B4
76:73	B5
77	P0
78	P1
79	P2
80	P3
81	P4
82	P5
127:83	indices

Mode 1: Color only, 2 lines (TWO), 6-bit endpoints with one shared P-bit per line, 3-bit indices, 64 partitions

Bit	Description
1:0	mode = 10
7:2	partition
13:8	R0
19:14	R1
25:20	R2
31:26	R3
37:32	G0

Bit	Description
43:38	G1
49:44	G2
55:50	G3
61:56	B0
67:62	B1
73:68	B2
79:74	B3
80	P0
81	P1
127:82	indices

Mode 2: Color only, 3 lines (THREE), 5-bit endpoints, 2-bit indices, 64 partitions

Bit	Description
2:0	mode = 100
8:3	partition
13:9	R0
18:14	R1
23:19	R2
28:24	R3
33:29	R4
38:34	R5
43:39	G0
48:44	G1
53:49	G2
58:54	G3
63:59	G4
68:64	G5
73:69	B0
78:74	B1
83:79	B2
88:84	B3
93:89	B4
98:94	B5
127:99	indices



Mode 3: Color only, 2 lines (TWO), 7-bit endpoints with one P-bit per endpoint, 2-bit indices, 64 partitions

Bit	Description
3:0	mode = 1000
9:4	partition
16:10	R0
23:17	R1
30:24	R2
37:31	R3
44:38	G0
51:45	G1
58:52	G2
65:59	G3
72:66	B0
79:73	B1
86:80	B2
93:87	B3
94	P0
95	P1
96	P2
97	P3
127:98	indices

Mode 4: Color and alpha, 1 line (ONE), 5-bit color endpoints, 6-bit alpha endpoints, 16 2-bit indices, 16 3-bit indices, 2-bit component rotation, 1-bit index selector

Bit	Description
4:0	mode = 10000
6:5	rotation
7	index selector
12:8	R0
17:13	R1
22:18	G0
27:23	G1
32:28	B0
37:33	B1
43:38	A0
49:44	A1
80:50	2-bit indices

Bit	Description
127:81	3-bit indices

Mode 5: Color and alpha, 1 line (ONE), 7-bit color endpoints, 8-bit alpha endpoints, 2-bit color indices, 2-bit alpha indices, 2-bit component rotation

Bit	Description
5:0	mode = 100000
7:6	rotation
14:8	R0
21:15	R1
28:22	G0
35:29	G1
42:36	B0
49:43	B1
57:50	A0
65:58	A1
96:66	color indices
127:97	alpha indices

Mode 6: Combined color and alpha, 1 line (ONE), 7-bit endpoints with one P-bit per endpoint, 4-bit indices

Bit	Description
6:0	mode = 1000000
13:7	R0
20:14	R1
27:21	G0
34:28	G1
41:35	B0
48:42	B1
55:49	A0
62:56	A1
63	P0
64	P1
127:65	indices



Mode 7: Combined color and alpha, 2 lines (TWO), 5-bit endpoints with one P-bit per endpoint, 2-bit indices, 64 partitions

Bit	Description
7:0	mode = 10000000
13:8	partition
18:14	R0
23:19	R1
28:24	R2
33:29	R3
38:34	G0
43:39	G1
48:44	G2
53:49	G3
58:54	B0
63:59	B1
68:64	B2
73:69	B3
78:74	A0
83:79	A1
88:84	A2
93:89	A3
94	P0
95	P1
96	P2
97	P3
127:98	indices

Undefined mode values (bits 7:0 = 00000000) return zero in the RGB channels.

The indices fields are variable in length and due to the different locations of the fix-up indices depending on partition set there are a very large number of possible configurations. Each mode above indicates how many bits each index has, and the fix-up indices (one in ONE mode, two in TWO mode, and three in THREE mode) each have one less bit than indicated. However, the indices are always packed into the index fields according to the table below, with the specific bit assignments of each texel following the rules just given.

Bit	Description
LSBs	texel[0][0] index
	texel[0][1] index
	texel[0][2] index

Bit	Description
	texel[0][3] index
	texel[1][0] index
	texel[1][1] index
	texel[1][2] index
	texel[1][3] index
	texel[2][0] index
	texel[2][1] index
	texel[2][2] index
	texel[2][3] index
	texel[3][0] index
	texel[3][1] index
	texel[3][2] index
MSBs	texel[3][3] index

Endpoint Computation

The endpoints can be defined with different precision depending on mode, as shown above. This section describes how the endpoints are computed from the bits in the compression block. Each component of each endpoint follows the same steps.

If a P-bit is defined for the endpoint, it is first added as an additional LSB at the bottom of the endpoint value. The endpoint is then bit-replicated to create an 8-bit fixed point endpoint value with a range from 0x00 to 0xFF.

Palette Color Computation

The next step involves computing the color palette values that provide the available values for each texel's color. The color palette for each line consists of the two endpoint colors plus 2, 6, or 14 interpolated colors, depending on the number of bits in the indices. Again each channel is processed independently.

The equation to compute each palette color with index i , given two endpoints is as follows, using the tables below to determine the weight for each palette index:

$$\text{palette}[i] = (E0 * (64 - \text{weight}[i]) + E1 * \text{weight}[i] + 32) \gg 6$$

2-bit index weights:

palette index	0	1	2	3
weight	0	21	43	64

3-bit index weights:

palette index	0	1	2	3	4	5	6	7
weight	0	9	18	27	37	46	55	64



4-bit index weights:

palette index	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
weight	0	4	9	13	17	21	26	30	34	38	43	47	51	55	60	64

The two end palette indices are equal to the two endpoints given that the weights are 0 and 64. In the above equation E0 and E1 represent the even-numbered and odd-numbered endpoints computed in the previous step for the component and line currently being computed.

Texel Selection

The final step is to select the appropriate palette index for each texel. This index then selects the 8-bit per channel palette value, which is interpreted as an 8-bit UNORM value for input into the filter (In BC7_UNORM_SRGB to UNORM values first go through inverse gamma conversion). This procedure differs depending on whether the mode is ONE, TWO, or THREE.

ONE Mode

In ONE mode, there is only one set of palette colors, thus there is only a single "partition set" defined, with all texels selecting line 0 and texel [0][0] being the "fix-up index" with one less bit in the index.

TWO Mode

64 partitions are defined for TWO, which are defined below. Each of the 64 cases shows the 4x4 block of texels, and is indexed by adding its hexadecimal row number (00-3C) to its column number (0-3). Each texel in the 4x4 block indicates whether that texel is part of line 0 (endpoints 0 and 1) or line 1 (endpoints 2 and 3). Each case has one texel each of "[0]" and "[1]", the index that this is at is termed the "fix-up index". These texels have one less bit in the index.

	0				1				2				3			
00	[0]	0	1	1	[0]	0	0	1	[0]	1	1	1	[0]	0	0	1
	0	0	1	1	0	0	0	1	0	1	1	1	0	0	1	1
	0	0	1	1	0	0	0	1	0	1	1	1	0	0	1	1
	0	0	1	[1]	0	0	0	[1]	0	1	1	[1]	0	1	1	[1]
04	[0]	0	0	0	[0]	0	1	1	[0]	0	0	1	[0]	0	0	0
	0	0	0	1	0	1	1	1	0	0	1	1	0	0	0	1
	0	0	0	1	0	1	1	1	0	1	1	1	0	0	1	1
	0	0	1	[1]	1	1	1	[1]	1	1	1	[1]	0	1	1	[1]
08	[0]	0	0	0	[0]	0	1	1	[0]	0	0	0	[0]	0	0	0
	0	0	0	0	0	1	1	1	0	0	0	1	0	0	0	0
	0	0	0	1	1	1	1	1	0	1	1	1	0	0	0	1
	0	0	1	[1]	1	1	1	[1]	1	1	1	[1]	0	1	1	[1]
0C	[0]	0	0	1	[0]	0	0	0	[0]	0	0	0	[0]	0	0	0
	0	1	1	1	0	0	0	0	1	1	1	1	0	0	0	0

	0				1				2				3			
	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0
	1	1	1	[1]	1	1	1	[1]	1	1	1	[1]	1	1	1	[1]
10	[0]	0	0	0	[0]	1	[1]	1	[0]	0	0	0	[0]	1	[1]	1
	1	0	0	0	0	0	0	1	0	0	0	0	0	0	1	1
	1	1	1	0	0	0	0	0	[1]	0	0	0	0	0	0	1
	1	1	1	[1]	0	0	0	0	1	1	1	0	0	0	0	0
14	[0]	0	[1]	1	[0]	0	0	0	[0]	0	0	0	[0]	1	1	1
	0	0	0	1	1	0	0	0	0	0	0	0	0	0	1	1
	0	0	0	0	[1]	1	0	0	[1]	0	0	0	0	0	1	1
	0	0	0	0	1	1	1	0	1	1	0	0	0	0	0	[1]
18	[0]	0	[1]	1	[0]	0	0	0	[0]	1	[1]	0	[0]	0	[1]	1
	0	0	0	1	1	0	0	0	0	1	1	0	0	1	1	0
	0	0	0	1	[1]	0	0	0	0	1	1	0	0	1	1	0
	0	0	0	0	1	1	0	0	0	1	1	0	1	1	0	0
1C	[0]	0	0	1	[0]	0	0	0	[0]	1	[1]	1	[0]	0	[1]	1
	0	1	1	1	1	1	1	1	0	0	0	1	1	0	0	1
	[1]	1	1	0	[1]	1	1	1	1	0	0	0	1	0	0	1
	1	0	0	0	0	0	0	0	1	1	1	0	1	1	0	0
20	[0]	1	0	1	[0]	0	0	0	[0]	1	0	1	[0]	0	1	1
	0	1	0	1	1	1	1	1	1	0	[1]	0	0	0	1	1
	0	1	0	1	0	0	0	0	0	1	0	1	[1]	1	0	0
	0	1	0	[1]	1	1	1	[1]	1	0	1	0	1	1	0	0
24	[0]	0	[1]	1	[0]	1	0	1	[0]	1	1	0	[0]	1	0	1
	1	1	0	0	0	1	0	1	1	0	0	1	1	0	1	0
	0	0	1	1	[1]	0	1	0	0	1	1	0	1	0	1	0
	1	1	0	0	1	0	1	0	1	0	0	[1]	0	1	0	[1]
28	[0]	1	[1]	1	[0]	0	0	1	[0]	0	[1]	1	[0]	0	[1]	1
	0	0	1	1	0	0	1	1	0	0	1	0	1	0	1	1
	1	1	0	0	[1]	1	0	0	0	1	0	0	1	1	0	1
	1	1	1	0	1	0	0	0	1	1	0	0	1	1	0	0
2C	[0]	1	[1]	0	[0]	0	1	1	[0]	1	1	0	[0]	0	0	0
	1	0	0	1	1	1	0	0	0	1	1	0	0	1	[1]	0
	1	0	0	1	1	1	0	0	1	0	0	1	0	1	1	0
	0	1	1	0	0	0	1	[1]	1	0	0	[1]	0	0	0	0
30	[0]	1	0	0	[0]	0	[1]	0	[0]	0	0	0	[0]	0	0	0
	1	1	[1]	0	0	1	1	1	0	0	[1]	0	0	1	0	0
	0	1	0	0	0	0	1	0	0	1	1	1	[1]	1	1	0
	0	0	0	0	0	0	0	0	0	0	1	0	0	1	0	0



	0				1				2				3			
34	[0]	1	1	0	[0]	0	1	1	[0]	1	[1]	0	[0]	0	[1]	1
	1	1	0	0	0	1	1	0	0	0	1	1	1	0	0	1
	1	0	0	1	1	1	0	0	1	0	0	1	1	1	0	0
	0	0	1	[1]	1	0	0	[1]	1	1	0	0	0	1	1	0
38	[0]	1	1	0	[0]	1	1	0	[0]	1	1	1	[0]	0	0	1
	1	1	0	0	0	0	1	1	1	1	1	0	1	0	0	0
	1	1	0	0	0	0	1	1	1	0	0	0	1	1	1	0
	1	0	0	[1]	1	0	0	[1]	0	0	0	[1]	0	1	1	[1]
3C	[0]	0	0	0	[0]	0	[1]	1	[0]	0	[1]	0	[0]	1	0	0
	1	1	1	1	0	0	1	1	0	0	1	0	0	1	0	0
	0	0	1	1	1	1	1	1	1	1	1	0	0	1	1	1
	0	0	1	[1]	0	0	0	0	1	1	1	0	0	1	1	[1]

THREE Mode

64 partitions are defined for THREE, which are defined below. Each of the 64 cases shows the 4x4 block of texels, and is indexed by adding its hexadecimal row number (00-3C) to its column number (0-3). Each texel in the 4x4 block indicates whether that texel is part of line 0 (endpoints 0 and 1), line 1 (endpoints 2 and 3), or line 2 (endpoints 4 and 5). Each case has one texel each of "[0]", "[1]", and "[2]", the index that this is at is termed the "fix-up index". These texels have one less bit in the index.

	0				1				2				3			
00	[0]	0	1	[1]	[0]	0	0	[1]	[0]	0	0	0	[0]	2	2	[2]
	0	0	1	1	0	0	1	1	2	0	0	1	0	0	2	2
	0	2	2	1	[2]	2	1	1	[2]	2	1	1	0	0	1	1
	2	2	2	[2]	2	2	2	1	2	2	1	[1]	0	1	1	[1]
04	[0]	0	0	0	[0]	0	1	[1]	[0]	0	2	[2]	[0]	0	1	1
	0	0	0	0	0	0	1	1	0	0	2	2	0	0	1	1
	[1]	1	2	2	0	0	2	2	1	1	1	1	[2]	2	1	1
	1	1	2	[2]	0	0	2	[2]	1	1	1	[1]	2	2	1	[1]
08	[0]	0	0	0	[0]	0	0	0	[0]	0	0	0	[0]	0	1	2
	0	0	0	0	1	1	1	1	1	1	[1]	1	0	0	[1]	2
	[1]	1	1	1	[1]	1	1	1	2	2	2	2	0	0	1	2
	2	2	2	[2]	2	2	2	[2]	2	2	2	[2]	0	0	1	[2]
0C	[0]	1	1	2	[0]	1	2	2	[0]	0	1	[1]	[0]	0	1	[1]
	0	1	[1]	2	0	[1]	2	2	0	1	1	2	2	0	0	1
	0	1	1	2	0	1	2	2	1	1	2	2	[2]	2	0	0
	0	1	1	[2]	0	1	2	[2]	1	2	2	[2]	2	2	2	0

	0				1				2				3			
10	[0]	0	0	[1]	[0]	1	1	[1]	[0]	0	0	0	[0]	0	2	[2]
	0	0	1	1	0	0	1	1	1	1	2	2	0	0	2	2
	0	1	1	2	[2]	0	0	1	[1]	1	2	2	0	0	2	2
	1	1	2	[2]	2	2	0	0	1	1	2	[2]	1	1	1	[1]
14	[0]	1	1	[1]	[0]	0	0	[1]	[0]	0	0	0	[0]	0	0	0
	0	1	1	1	0	0	0	1	0	0	[1]	1	1	1	0	0
	0	2	2	2	[2]	2	2	1	0	1	2	2	[2]	2	[1]	0
	0	2	2	[2]	2	2	2	1	0	1	2	[2]	2	2	1	0
18	[0]	1	2	[2]	[0]	0	1	2	[0]	1	1	0	[0]	0	0	0
	0	[1]	2	2	0	0	1	2	1	2	[2]	1	0	1	[1]	0
	0	0	1	1	[1]	1	2	2	[1]	2	2	1	1	2	[2]	1
	0	0	0	0	2	2	2	[2]	0	1	1	0	1	2	2	1
1C	[0]	0	2	2	[0]	1	1	0	[0]	0	1	1	[0]	0	0	0
	1	1	0	2	0	[1]	1	0	0	1	2	2	2	0	0	0
	[1]	1	0	2	2	0	0	2	0	1	[2]	2	[2]	2	1	1
	0	0	2	[2]	2	2	2	[2]	0	0	1	[1]	2	2	2	[1]
20	[0]	0	0	0	[0]	2	2	[2]	[0]	0	1	[1]	[0]	1	2	0
	0	0	0	2	0	0	2	2	0	0	1	2	0	[1]	2	0
	[1]	1	2	2	0	0	1	2	0	0	2	2	0	1	[2]	0
	1	2	2	[2]	0	0	1	[1]	0	2	2	[2]	0	1	2	0
24	[0]	0	0	0	[0]	1	2	0	[0]	1	2	0	[0]	0	1	1
	1	1	[1]	1	1	2	0	1	2	0	1	2	2	2	0	0
	2	2	[2]	2	[2]	0	[1]	2	[1]	[2]	0	1	1	1	[2]	2
	0	0	0	0	0	1	2	0	0	1	2	0	0	0	1	[1]
28	[0]	0	1	1	[0]	1	0	[1]	[0]	0	0	0	[0]	0	2	2
	1	1	[2]	2	0	1	0	1	0	0	0	0	1	[1]	2	2
	2	2	0	0	2	2	2	2	[2]	1	2	1	0	0	2	2
	0	0	1	[1]	2	2	2	[2]	2	1	2	[1]	1	1	2	[2]
2C	[0]	0	2	[2]	[0]	2	2	0	[0]	1	0	1	[0]	0	0	0
	0	0	1	1	1	2	[2]	1	2	2	[2]	2	2	1	2	1
	0	0	2	2	0	2	2	0	2	2	2	2	[2]	1	2	1
	0	0	1	[1]	1	2	2	[1]	0	1	0	[1]	2	1	2	[1]
30	[0]	1	0	[1]	[0]	2	2	[2]	[0]	0	0	2	[0]	0	0	0
	0	1	0	1	0	1	1	1	1	[1]	1	2	2	[1]	1	2
	0	1	0	1	0	2	2	2	0	0	0	2	2	1	1	2
	2	2	2	[2]	0	1	1	[1]	1	1	1	[2]	2	1	1	[2]



	0				1				2				3			
34	[0]	2	2	2	[0]	0	0	2	[0]	1	1	0	[0]	0	0	0
	0	[1]	1	1	1	1	1	2	0	[1]	1	0	0	0	0	0
	0	1	1	1	[1]	1	1	2	0	1	1	0	2	1	[1]	2
	0	2	2	[2]	0	0	0	[2]	2	2	2	[2]	2	1	1	[2]
38	[0]	1	1	0	[0]	0	2	2	[0]	0	2	2	[0]	0	0	0
	0	[1]	1	0	0	0	1	1	1	1	2	2	0	0	0	0
	2	2	2	2	0	0	[1]	1	[1]	1	2	2	0	0	0	0
	2	2	2	[2]	0	0	2	[2]	0	0	2	[2]	2	[1]	1	[2]
3C	[0]	0	0	[2]	[0]	2	2	2	[0]	1	0	[1]	[0]	1	1	[1]
	0	0	0	1	1	2	2	2	2	2	2	2	2	0	1	1
	0	0	0	2	0	2	2	2	2	2	2	2	[2]	2	0	1
	0	0	0	[1]	[1]	2	2	[2]	2	2	2	[2]	2	2	2	0

Video Pixel/Texel Formats

This section describes the "video" pixel/texel formats with respect to memory layout. See the Overlay chapter for a description of how the Y, U, V components are sampled.

Packed Memory Organization

Color components are all 8 bits in size for YUV formats. For YUV 4:2:2 formats each DWord will contain two pixels and only the byte order affects the memory organization.

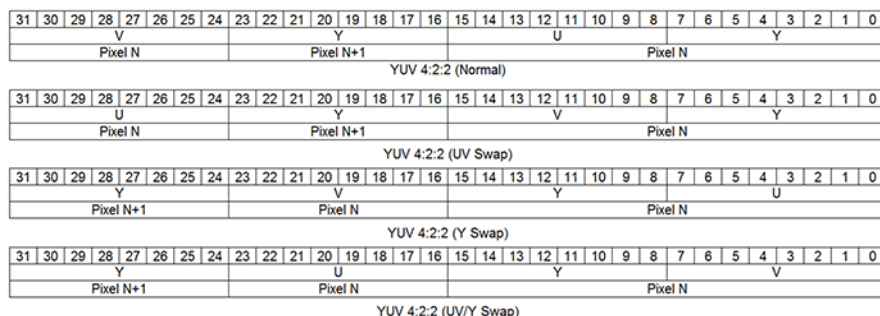
The following four YUV 4:2:2 surface formats are supported, listed with alternate names:

- YCRCB_NORMAL (YUYV/YUY2)
- YCRCB_SWAPUVY (VYUY) (R8G8_B8G8_UNORM)
- YCRCB_SWAPUV (YVYU) (G8R8_G8B8_UNORM)
- YCRCB_SWAPY (UYVY)

The channels are mapped as follows:

Cr (V)	Red
Y	Green
Cb (U)	Blue

Memory layout of packed YUV 4:2:2 formats



Planar Memory Organization

Planar formats use what could be thought of as separate buffers for the three color components. Because there is a separate stride for the Y and U/V data buffers, several memory footprints can be supported.

The 3D sampler supports direct sampling and filtering of planar video surfaces such as YV12 and NV12.

Programming Note	
Context:	Tiling of Planar Surface
<p>Tiling of planar surfaces (tileX, tileY, tileYf, or tileYs) is only supported for planar surfaces where the chroma plane is full-pitch (e.g. NV21). In this case, the field Y Offset for U or UV Plane in the RENDER_SURFACE_STATE must be programmed to force the UV plane to be at the start of a tile.</p>	

Programming Note	
Context:	YV12/YV21 Surface Pitch Restriction
<p>The Surface Pitch defined in RENDER_SURFACE_STATE must be a multiple of 64Bytes for YV12 and YV21 surfaces.</p>	


Programming Note	
Context:	NV21 Support
<p>Sampling of NV21 surface format is supported by swapping the U and V channels when sampling the surface. This can be done by programming the Shader Channel Select in the RENDER_SURFACE_STATE for the Red and Blue Channels.</p>	

The 3D sampler supports 12, and 16-bit planar video surface formats known collectively as P016. They are only supported for Sample_Unorm.

The following figure shows two types of memory organization for the YUV 4:2:0 planar video data:

1. The memory organization of the common YV12 data, where all three planes are contiguous and the strides of U and V components are half of that of the Y component.
2. An alternative memory structure that the addresses of the three planes are independent but satisfy certain alignment restrictions.

YUV 4:2:0 Format Memory Organization




B6684-01

The following figure shows memory organization of the planar YUV 4:1:0 format where the planes are contiguous.

Note: The chroma planes (U and V), when separate (case b above) are treated as half-pitch with respect to the Y plane.

Workaround
When using Planar formats for YUV with half-pitch chroma planes (e.g. YV12), and fenced tiling is not supported
LINEAR filtering of Planar YUV surfaces such as YV12 using the 3D sampler is done after the U and V have been replicated to form a YUV444 texels. This means that the U and V components will effectively be point-sampled rather than filtered. Achieve true filtering of the U and V components, the 3 planes of the YUV surface must be bound as separate surfaces, and the filtering must be done on each individually.
LINEAR filtering of Planar YUV surfaces such as YV12 using 3D sampler can be done across all planes, but only with replicated U and V components. Likewise, only replication is supported for Anisotropic and sample_min/sample_max. To achieve true linear, anisotropic or min/max filtering the 3 plans of the YUV surface must be bound as separate surfaces, and the filtering done on each individually.

YUV 4:1:0 Format Memory Organization



B6685-01

The table below shows how position within a Planar YUV surface chroma plane is calculated for various cases of U and V pitch and position. It also shows restrictions on the alignment of the chroma planes in memory for non-interleaved (YV12) and interleaved chroma (e.g. NV12) is used.

Case	Interleave Chroma	Pitch	Vertical U/V Offset
YUV with Half Pitch Chroma	No	Half	<u>When U is below Y</u> $Y_Uoffset = Y_Height * 2$ $Y_Voffset = Y_Height * 2 + V_Height$ <u>When V is below Y</u> $Y_Uoffset = Y_Height * 2 + V_Height$ $Y_Voffset = Y_Height * 2$
YUV with Full Pitch Chroma	Yes	Full	<u>When U is below Y</u> $Y_Uoffset = Y_Height$ $Y_Voffset = Y_Height + V_Height$ <u>When V is below Y</u> $Y_Uoffset = Y_Height + V_Height$ $Y_Voffset = Y_Height$
YUV for Media Sampling	Yes	Always Full	Same as 3D full pitch

Programming Note

Context:

Planar YUV surfaces cannot be 1D surface types.

Because there is a requirement that the height of the Y plane of a planar surface must be a greater than 1, it cannot be programmed to be a **Surface Type** of SURFTYPE_1D.



Programming Note	
Context:	MIP Filtering
Surface state cannot have (MIP Mode Filter != NONE) for Planar YUV surfaces (e.g. PLANAR_420_8).	

Programming Note	
Context:	Standard Tiling
Planar YUV does not support MIP Tails as part of Standard Tiling. The MIP Tail Start field in RENDER_SURFACE_STATE must be programmed to 15.	

Programming Note	
Context:	Quilted and Planar
Planar YUV is not supported for Quilted surfaces.	

Programming Note	
Context:	
Planar YUV is not supported with Corner Texel Mode	

Additional Video Formats

Dx YUV format name	DxRGB format name	GPGPU Image Type	GPGPU Image Data Type	Packed / Planar	# Surface States	Surface Format #1	Surface Format #2	Surface Format #3	Support By
DXGI_FORMAT_AYUV	DXGI_FORMAT_R8G8B8A8_UNORM (V->R8, U->G8, Y->B8, A->A8)			Packed	1	R8G8B8A8_UNORM	NA	NA	Sampler, PB
DXGI_FORMAT_AYUV	DXGI_FORMAT_R8G8B8A8_UINT (V->R8, U->G8, Y->B8, A->A8)			Packed	1	R8G8B8A8_UINT	NA	NA	Sampler, HDC, PB
DXGI_FORMAT_YUY2	DXGI_FORMAT_R8G8B8A8_UNORM (Y0->R8, U0->G8, Y1->B8, V0->A8)			Packed	1	R8G8B8A8_UNORM	NA	NA	Sampler,
NA	NA	CL_YCbYCr	CL_UNORM_INT16	Packed	1	R16G16B16A16_UINT	NA	NA	Sampler, HDC

Dx YUV format name	DxRGB format name	GPGPU Image Type	GPGPU Image Data Type	Packed / Planar	# Surface States	Surface Format #1	Surface Format #2	Surface Format #3	Support By
		>R16, U0->G16, Y1->B16, V0->A16							
NA	NA	CL_YCbYCr (Y0->R16, U0->G16, Y1->B16, V0->A16	CL_UNORM_INT12	Packed	1	R16G16B16A16_UINT	NA	NA	Sampler, HDC
NA	NA	CL_YCbYCr (Y0->R16, U0->G16, Y1->B16, V0->A16	CL_UNORM_INT10	Packed	1	R16G16B16A16_UINT	NA	NA	Sampler, HDC
DXGI_FORMAT_YUY2	DXGI_FORMAT_R8G8B8A8_UINT (Y0->R8, U0->G8, Y1->B8, V0->A8)	CL_YCbYCr (Y0->R8, U0->G8, Y1->B8, V0->A8	CL_UNORM_INT8	Packed	1	R8G8B8A8_UINT	NA	NA	Sampler, HDC
DXGI_FORMAT_YUY2	DXGI_FORMAT_R8G8B8G8_UNORM			Packed	1	R8G8B8A8_UNORM In this case the width of the view will appear to be twice the R8G8B8A8 view, with hardware reconstruction of RGBA done automatically on read (and before filtering).	NA	NA	Sampler

Dx YUV format name	DxRGB format name	GPGPU Image Type	GPGPU Image Data Type	Packed / Planar	# Surface States	Surface Format #1	Surface Format #2	Surface Format #3	Support By
NA	NA	CL_CbYCrY (U0->R16, Y0->G16, V0->B16, Y1->A16	CL_UNORM_INT16	Packed	1	R16G16B16A16_UINT	NA	NA	Sampler, HDC
NA	NA	CL_CbYCrY (U0->R16, Y0->G16, V0->B16, Y1->A16	CL_UNORM_INT12	Packed	1	R16G16B16A16_UINT	NA	NA	Sampler, HDC
NA	NA	CL_CbYCrY (U0->R16, Y0->G16, V0->B16, Y1->A16	CL_UNORM_INT10	Packed	1	R16G16B16A16_UINT	NA	NA	Sampler, HDC
NA	NA	CL_CbYCrY (U0->R8, Y0->G8, V0->B8, Y1->A8	CL_UNORM_INT8	Packed	1	R8G8B8A8_UINT	NA	NA	Sampler, HDC
DXGI_FORMAT_NV12	Y = DXGI_FORMAT_R8_UNORM U/V = DXGI_FORMAT_R8G8_UNORM (U->R8, V->G8)			Planar	2	R8_UNORM	R8G8_UNORM chroma pixel dimensions 1/2 in both x and y from the Luma	NA	Sampler, PB

Dx YUV format name	DxRGB format name	GPGPU Image Type	GPGPU Image Data Type	Packed / Planar	# Surface States	Surface Format #1	Surface Format #2	Surface Format #3	Support By
							view		
NA	NA	CL_Y_CrCb (Y -> R16) (U->R16) (V->R16)	CL_UNORM_INT16	Planar	3	R16_UNIT	R16_UNIT	R16_UNIT	Sampler
NA	NA	CL_Y_CrCb (Y -> R16) (U->R16) (V->R16)	CL_UNORM_INT12	Planar	3	R16_UNIT	R16_UNIT	R16_UNIT	Sampler
NA	NA	CL_Y_CrCb (Y -> R16) (U->R16) (V->R16)	CL_UNORM_INT10	Planar	3	R16_UNIT	R16_UNIT	R16_UNIT	Sampler
NA	NA	CL_Y_CrCb (Y -> R8) (U->R8) (V->R8)	CL_UNORM_INT8	Planar	3	R8_UNIT	R8_UNIT	R8_UNIT	Sampler
NA	NA	CL_Y_CrCb (Y -> R16) (U->R16, V->G16)	CL_UNORM_INT16	Planar	2	R16_UNIT	R16G16_UINT chroma pixel dimensions halved in both x and y from the Luma view	NA	Sampler, HDC

Dx YUV format name	DxRGB format name	GPGPU Image Type	GPGPU Image Data Type	Packed / Planar	# Surface States	Surface Format #1	Surface Format #2	Surface Format #3	Support By
NA	NA	CL_Y_CrCb (Y -> R16) (U->R16, V->G16)	CL_UNORM_INT12	Planar	2	R16_UNIT	R16G16_UINT chomra pixel dimensions halved in both x and y from the Luma view	NA	Sampler, HDC
NA	NA	CL_Y_CrCb (Y -> R16) (U->R16, V->G16)	CL_UNORM_INT10	Planar	2	R16_UNIT	R16G16_UINT chomra pixel dimensions halved in both x and y from the Luma view	NA	Sampler, HDC
DXGI_FORMAT_NV12	Y = DXGI_FORMAT_R8_UINT U/V = DXGI_FORMAT_R8G8_UINT (U->R8, V->G8)	CL_Y_CrCb (Y -> R8) (U->R8, V->G8)	CL_UNORM_INT8	Planar	2	R8_UNIT	R8G8_UINT chomra pixel dimensions halved in both x and y from the Luma view	NA	Sampler, HDC, PB
DXGI_FORMAT_NV11	Y = DXGI_FORMAT_R8_UNORM U/V = DXGI_FORMAT_R8G8_UNORM (U->R8, V->G8)			Planar	2	R8_UNORM	R8G8_UNORM chomra pixel dimensions 1/4 in both x and y	NA	Sampler, PB

Dx YUV format name	DxRGB format name	GPGPU Image Type	GPGPU Image Data Type	Packed / Planar	# Surface States	Surface Format #1	Surface Format #2	Surface Format #3	Support By
							from the Luma view		
DXGI_FORMAT_NV11	Y = DXGI_FORMAT_R8_UINT U/V = DXGI_FORMAT_R8G8_UINT (U->R8, V->G8)			Planar	2	R8_UNIT	R8G8_UINT chroma pixel dimensions 1/4 in both x and y from the Luma view	NA	Sampler, HDC, PB
DXGI_FORMAT_P016	Y = DXGI_FORMAT_R16_UNORM U/V = DXGI_FORMAT_R16G16_UNORM (U->R16, V->G16)			Planar	2	R16_UNORM	R16G16_UNORM chroma pixel dimensions 1/2 in both x and y from the Luma view	NA	Sampler, HDC, PB
DXGI_FORMAT_P016	Y = DXGI_FORMAT_R16_UINT U/V = DXGI_FORMAT_R16G16_UINT (U->R16, V->G16)			Planar	2	R16_UNIT	R16G16_UINT chroma pixel dimensions 1/2 in both x and y from the Luma view	NA	Sampler, PB
DXGI_FORMAT_P010	Y = DXGI_FORMAT_R16_UNORM U/V =			Planar	2	R16_UNORM	R16G16_UNORM chroma pixel	NA	Sampler, HDC, PB

Dx YUV format name	DxRGB format name	GPGPU Image Type	GPGPU Image Data Type	Packed / Planar	# Surface States	Surface Format #1	Surface Format #2	Surface Format #3	Support By
	DXGI_FORMAT_R16G16_UNORM (U->R16, V->G16)						dimensions 1/2 in both x and y from the Luma view		
DXGI_FORMAT_P010	Y = DXGI_FORMAT_R16_UINT U/V = DXGI_FORMAT_R16G16_UINT (U->R16, V->G16)			Planar	2	R16_UNIT	R16G16_UINT chroma pixel dimensions 1/2 in both x and y from the Luma view	NA	Sampler, HDC, PB
DXGI_FORMAT_Y216	DXGI_FORMAT_R16G16B16A16_UNORM (Y0->R16, U->G16, Y1->B16, V->A16).			Packed	1	R16G16B16A16_UNORM	NA	NA	Sampler,
DXGI_FORMAT_Y216	DXGI_FORMAT_R16G16B16A16_UINT (Y0->R16, U->G16, Y1->B16, V->A16).			Packed	1	R16G16B16A16_UINT	NA	NA	Sampler, HDC
DXGI_FORMAT_Y210	DXGI_FORMAT_R16G16B16A16_UNORM (Y0->R16, U->G16, Y1->B16, V->A16).			Packed	1	R16G16B16A16_UNORM	NA	NA	Sampler,
DXGI_FORMAT_Y210	DXGI_FORMAT_R16G16B16A16_UINT (Y0->R16, U->G16, Y1->B16, V->A16).			Packed	1	R16G16B16A16_UINT	NA	NA	Sampler, HDC
DXGI_FORMAT_Y416	DXGI_FORMAT_R16G16B16A16_UNORM			Packed	1	R16G16B16A16_UNORM	NA	NA	Sampler, HDC

Dx YUV format name	DxRGB format name	GPGPU Image Type	GPGPU Image Data Type	Packed / Planar	# Surface States	Surface Format #1	Surface Format #2	Surface Format #3	Support By
	(U->R16, Y->G16, V->B16, A->A16)								
DXGI_FORMAT_Y416	DXGI_FORMAT_R16G16B16A16_UINT (U->R16, Y->G16, V->B16, A->A16)			Packed	1	R16G16B16A16_UINT	NA	NA	Sampler,
DXGI_FORMAT_Y410	DXGI_FORMAT_R16G16B16A16_UNORM (U->R16, Y->G16, V->B16, A->A16)			Packed	1	R16G16B16A16_UNORM	NA	NA	Sampler, HDC
DXGI_FORMAT_Y410	DXGI_FORMAT_R16G16B16A16_UINT (U->R16, Y->G16, V->B16, A->A16)			Packed	1	R16G16B16A16_UINT	NA	NA	Sampler,
NA	NA	CL_CbCr (U->R16, V->G16)	CL_UNORM_INT16	Packed	1	R16G16_UINT	NA	NA	Sampler, HDC
NA	NA	CL_CbCr (U->R16, V->G16)	CL_UNORM_INT12	Packed	1	R16G16_UINT	NA	NA	Sampler, HDC
NA	NA	CL_CbCr (U->R16, V->G16)	CL_UNORM_INT10	Packed	1	R16G16_UINT	NA	NA	Sampler, HDC
NA	NA	CL_CbCr (U->R8, V->G8)	CL_UNORM_INT8	Packed	1	R8G8_UINT	NA	NA	Sampler, HDC
NA	NA	CL_Cb_Cr (U->R16) (V->R16)	CL_UNORM_INT16	Planar	2	R16_UNIT	R16_UNIT	NA	Sampler, HDC
NA	NA	CL_Cb_Cr	CL_UNORM_INT12	Planar	2	R16_UNIT	R16_UNIT	NA	Sampler, HDC

Dx YUV format name	DxRGB format name	GPGPU Image Type	GPGPU Image Data Type	Packed / Planar	# Surface States	Surface Format #1	Surface Format #2	Surface Format #3	Support By
		(U->R16) (V->R16)							
NA	NA	CL_Cb_Cr (U->R16) (V->R16)	CL_UNORM_INT10	Planar	2	R16_UNIT	R16_UNIT	NA	Sampler, HDC
NA	NA	CL_Cb_Cr (U->R8) (V->R8)	CL_UNORM_INT8	Planar	2	R8_UNIT	R8_UNIT	NA	Sampler, HDC
NA	NA	CL_Y (Y->R16)	CL_UNORM_INT16	Packed	1	R16_UNIT	NA	NA	Sampler, HDC
NA	NA	CL_Y (Y->R16)	CL_UNORM_INT12	Packed	1	R16_UNIT	NA	NA	Sampler, HDC
NA	NA	CL_Y (Y->R16)	CL_UNORM_INT10	Packed	1	R16_UNIT	NA	NA	Sampler, HDC
NA	NA	CL_Y (Y->R8)	CL_UNORM_INT8	Packed	1	R8_UNIT	NA	NA	Sampler, HDC
NA	NA	CL_Cb (V->R16)	CL_UNORM_INT16	Packed	1	R16_UNIT	NA	NA	Sampler, HDC
NA	NA	CL_Cb (V->R16)	CL_UNORM_INT12	Packed	1	R16_UNIT	NA	NA	Sampler, HDC
NA	NA	CL_Cb (V->R16)	CL_UNORM_INT10	Packed	1	R16_UNIT	NA	NA	Sampler, HDC
NA	NA	CL_Cb (V->R8)	CL_UNORM_INT8	Packed	1	R8_UNIT	NA	NA	Sampler, HDC
NA	NA	CL_Cb	CL_UNORM_INT16	Packed	1	R16_UNIT	NA	NA	Sampler, HDC

Dx YUV format name	DxRGB format name	GPGPU Image Type	GPGPU Image Data Type	Packed / Planar	# Surface States	Surface Format #1	Surface Format #2	Surface Format #3	Support By
		(U->R16)							
NA	NA	CL_Cb (U->R16)	CL_UNORM_INT12	Packed	1	R16_UNIT	NA	NA	Sampler, HDC
NA	NA	CL_Cb (U->R16)	CL_UNORM_INT10	Packed	1	R16_UNIT	NA	NA	Sampler, HDC
NA	NA	CL_Cb (U->R8)	CL_UNORM_INT8	Packed	1	R8_UNIT	NA	NA	Sampler, HDC



Raw Format

A format called "RAW" is available that is only supported with the untyped surface read/write, block, scattered, and atomic operation data port messages. It means that the surface has no inherent format. Surfaces of type RAW are addressed with byte-based offsets. The RAW surface format can be applied only to surface types of BUFFER and SCRATCH.

Surface Memory Organizations

See *Memory Interface Functions* chapter for a discussion of tiled vs. linear surface formats.

Display, Overlay, Cursor Surfaces

These surfaces are memory image buffers (planes) used to refresh a display device in non-VGA mode. See the Display chapter for specifics on how these surfaces are defined/used.

2D Render Surfaces

These surfaces are used as general source and/or destination operands in 2D BLT operations.

Note that there is no coherency between 2D render surfaces and the texture cache. Software must explicitly invalidate the texture cache before using a texture that has been modified via the BLT engine.

See the 2D Instruction and 2D Rendering chapters for specifics on how these surfaces are used, restrictions on their size, placement, etc.

2D Monochrome Source

These 1 BPP (bit per pixel) surfaces are used as source operands to certain 2D BLT operations, where the BLT engine expands the 1 BPP source to the required color depth.

The texture cache stores any monochrome sources. There is no mechanism to maintain coherency between 2D render surfaces and texture-cached monochrome sources. Software must explicitly invalidate the texture cache before using a memory-based monochrome source that has been modified via the BLT engine. (Here the assumption is that SW enforces memory-based monochrome source surfaces as read-only surfaces.)

See the 2D Instruction and 2D Rendering chapters for specifics on how these surfaces are used, restrictions on their size, placement, coherency rules, etc.

2D Color Pattern

Color pattern surfaces are used as special pattern operands in 2D BLT operations.

The device uses the texture cache to store color patterns. There is no mechanism to maintain coherency between 2D render surfaces and (texture)-cached color patterns. Software is required to explicitly invalidate the texture cache before using a memory-based color pattern that has been modified via the BLT engine. (Here the assumption is that SW enforces memory-based color pattern surfaces as read-only surfaces.)

See the *2D Instruction* and *2D Rendering* chapters for specifics on how these surfaces are used, restrictions on their size, placement, etc.

3D Color Buffer (Destination) Surfaces

3D Color Buffer surfaces hold per-pixel color values for use in the 3D Pipeline. The 3D Pipeline always requires a Color Buffer to be defined.

See the Non-Video Pixel/Texel Formats section in this chapter for details on the Color Buffer pixel formats. See the 3D Instruction and 3D Rendering chapters for Color Buffer usage details.

The Color Buffer is defined as the BUFFERID_COLOR_BACK memory buffer via the 3DSTATE_BUFFER_INFO instruction. That buffer can be mapped to LM or SM (snooped or unsnooped), and can be linear or tiled. When both the Depth and Color Buffers are tiled, the respective Tile Walk directions must match.

When a linear Color Buffer and a linear Depth Buffer are used together:

- The buffers may have different pitches, though both pitches must be a multiple of 32 bytes.
- The buffers must be co-aligned with a 32-byte region.

3D Depth Buffer Surfaces

Depth Buffer surfaces hold per-pixel depth values and per-pixel stencil values for use in the 3D Pipeline. The 3D Pipeline does not require a Depth Buffer in general, though a Depth Buffer is required to perform non-trivial Depth Test and Stencil Test operations.

The Depth Buffer is specified via the 3DSTATE_DEPTH_BUFFER command. See the description of that instruction in *Windower* for restrictions.

See *Depth Buffer Formats* below for a summary of the possible depth buffer formats. See the Depth Buffer Formats section in this chapter for details on the pixel formats. See the *Windower* and *DataPort* chapters for details on the usage of the Depth Buffer.

Depth Buffer Formats

DepthBufferFormat / DepthComponent	BPP (Bits Per Pixel)	Description
D32_FLOAT	32	32-bit floating point Z depth value
D24_UNORM_S8_UINT	32	24-bit fixed point Z depth value in lower 3 bytes, 8-bit stencil value in upper byte
D16_UNORM	16	16-bit fixed point Z depth value



3D Separate Stencil Buffer Surfaces

Separate Stencil Buffer surfaces hold per-pixel stencil values for use in the 3D Pipeline. Note that the 3D Pipeline does not require a Stencil Buffer to be allocated, though a Stencil Buffer is required to perform non-trivial Stencil Test operations.

Depth Buffer Formats summarizes Stencil Buffer formats. Refer to the Stencil Buffer Formats section in this chapter for details on the pixel formats. Refer to the *Windower* chapters for Stencil Buffer usage details.

The Stencil buffer is specified via the 3DSTATE_STENCIL_BUFFER command. See that instruction description in *Windower* for restrictions.

Depth Buffer Formats

DepthBufferFormat / DepthComponent	BPP (bits per pixel)	Description
R8_UNIT	8	8-bit stencil value in a byte

Surface Layout and Tiling

This section explains how various surface types (1D, 2D, 3D, and Cube) are laid out in memory. Most of the information in this section is independent of tiling. The concept of tiling can be laid on top of information. Wherever there is a specific difference it will be called out.

For Tiling (TileY, TileYs etc.), see the Address Tiling Function Introduction section which provides detailed information on how tiles are organized and laid out.

Maximum Surface Size in Bytes

In addition to restrictions on maximum height, width, and depth, surfaces are also restricted to a maximum size of 2^{44} bytes. All pixels within the surface must be contained within 2^{44} bytes of the base address.

NULL Page Support for Media Sampler

NULL support for VA Media sampler will process the returned data for NULL pages as if it was all zeros. Except for feature matching function, in which the NULL page indication will cause the distance function calculation to be ignored and a maximum distance value of 0xFF to be returned.

Tiling

To improve efficiency in memory accesses, most surfaces can be laid out using a tiling scheme.

Supported Legacy Tiling Modes:

- **TileX**

Additional Tiling Modes:


- **Tile4:** 4KB tiling
- **Tile64:** 64KB tiling mode

These modes are described in the Address Tiling Function Introduction section.

Typed Buffers

A typed buffer is an array of structures. Each structure contains up to 2048 bytes of elements. Each element is a single surface format using one of the supported surface formats depending on how the surface is being accessed. The surface pitch state for the surface specifies the size of each structure in bytes.

The buffer is stored in memory contiguously with each element in the structure packed together, and the first element in the next structure immediately following the last element of the previous structure. Buffers are supported only in linear memory.



B.6686-01

Typed buffers are accessed using a surface state for each structure element (a,b,c, etc. in the diagram above). The surface state for element "b" (for example) contains the surface format of element "b" (which may differ from other elements), the base address points to element "b" in the first structure (slice 0 of the array). The pitch for all of the elements in the buffer is the same value, and the surface type of each element is SURFTYPE_BUFFER.

The offset into the typed buffer is given by the following equation:

$$\text{Offset} = (V * \text{Pitch}) + U$$

MIP Layout

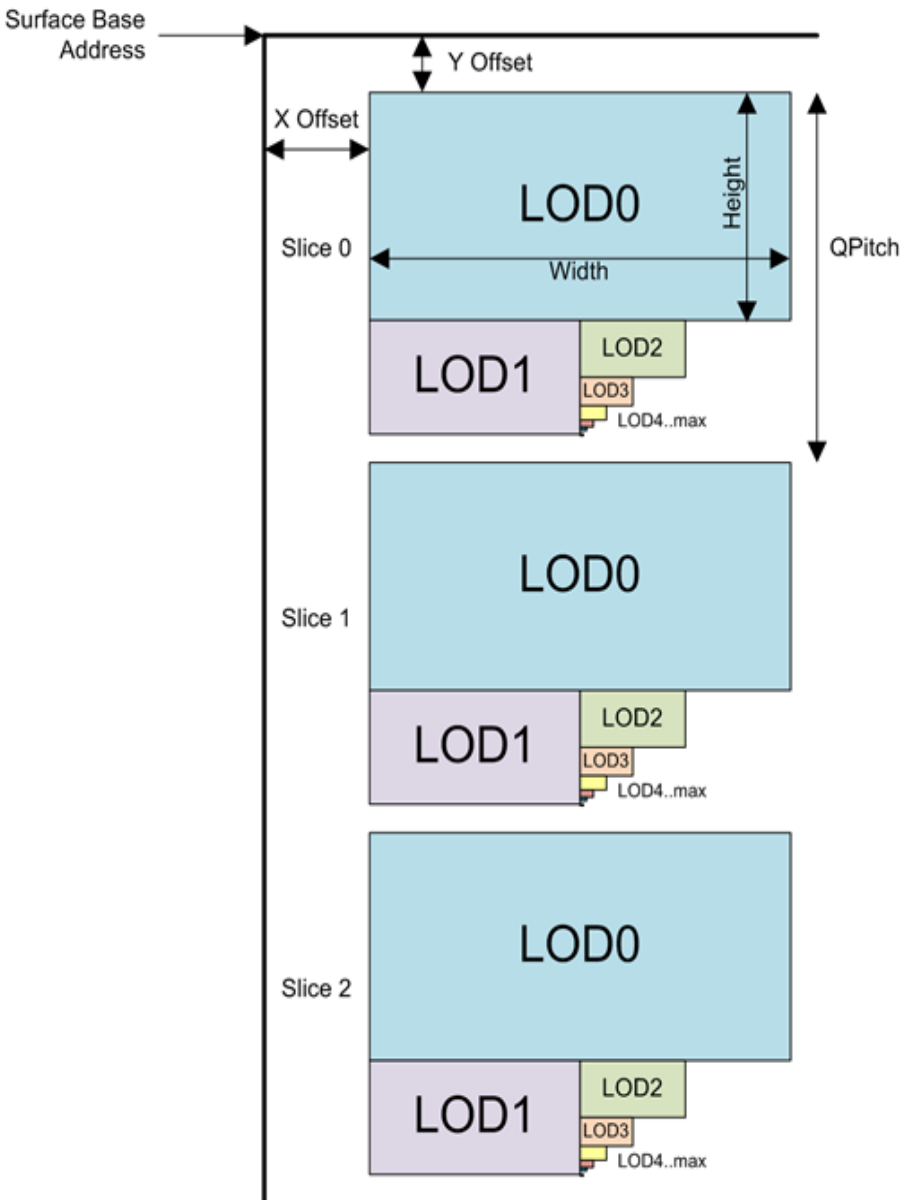
A surface can support multiple levels of details (LODs) or MIPs. The MIPCOUNT field in the RENDER_SURFACE_STATE defines how many MIPs a surface contains.

MIP0 or LOD0 is the largest, highest-detail MIP. The height, width and depth of this LOD is what is defined in the RENDER_SURFACE_STATE for that surface. Each subsequent

MIP is exactly one-half the height and width of the previous, making it 1/4th the size in memory.

The MIPs of a surface are laid out in memory using a 2-dimensional method as shown below. Volumetric and arrayed surfaces use multiple "slices" of this MIP layout, with each slice separated by QPITCH number of rows.

The diagram below shows many of the parameters of a 2D,2D Arrayed and 3D surface.



This 2-dimensional layout implies that there is padding required on the rows below LOD0 in order to ensure each row is the same number of texels.

If Tiling is enabled, then each MIP is laid out using one or more tiles. If Tile64 tiling is enabled, then some of the MIPs may actually be stored in a MIPTail which fits in a single 64K tile. The layout above, then only applied to MIPs which are not packed in the MIP Tail. Note that, depending on surface height the Vertical Alignment that surface can actually have the last few mips layed out below LOD1 for Tile4. Using MIP Tail (if enabled for Tile64) eliminates this possibility.

Raw (Untyped) Buffers

Raw buffers also use the surface type of SURFTYPE_BUFFER, but the surface format is RAW. These buffers are one-dimensional. They are accessed with a single U parameter which is a byte offset into the buffer. Raw buffers are also supported only in linear memory.

The offset into the raw buffer is given directly by the U parameter.

`Offset = U`

Scratch Buffers

A scratch buffer is a surface type that is based on the 2-dimensional Structured Buffer surface. It can be thought of as an array of buffers, where each buffer is a predefined number of DWords in size. The first coordinate (U) defines the array index, and the second coordinate (V) is a byte offset into the buffer which must be a multiple of 4 (DWord-aligned). A scratch buffer must be defined with **Surface Format RAW**.

A scratch buffer is used for temporary memory private to a thread. The physical thread ID is used as the array index to select the buffer, and the single coordinate is the byte offset into the private buffer.

The scratch buffer has only one dimension programmed in SURFACE_STATE which indicates the array size. The byte offset dimension (V) is assumed to be bounded only by the **Surface Pitch**.

The two-dimensional offset into the surface is defined directly by the U and V parameters. Scratch buffers are linear.

1D Surfaces

One-dimensional surfaces use a tiling mode of linear in RENDER_SURFACE_STATE. Therefore they do not support MIP Tails and use the **Horizontal Alignment** field in RENDER_SURFACE_STATE to determine spacing between MIPs. 1D surfaces are stored linearly in memory.

Programming Note	
Context:	Legacy 1D Tiling
<p>There is a legacy mode for representing a 1D surface as a 2D surface with a height of 1 texel. However, this mode is not recommended due to API compatibility. The Sampler Legacy 1D Map Layout Disable MMIO bit (bit 0, E194h) <i>must</i> be set to 1h to allow true 1D surfaces.</p>	


Programming Note

Context:

Legacy 1D Tiling

There is a legacy mode for representing a 1D surface as a 2D surface with a height of 1 texel. However, this mode is not recommended due to API compatibility. The MMIO bit (bit 0, E194h) *must* be set to 1h to allow true 1D surfaces.

Linear 1D surfaces are stored in a one-dimensional view of memory as follows:



Surface Pitch is ignored for 1D surfaces. **Surface QPitch** specifies the distance in pixels between array slices. QPitch should allow at least enough space for any mips that may be present.

A number of parameters are useful to determine where given pixels will be located on the 1D surface. First, the width for each LOD "L" is computed:

$$W_L = ((width > L) > 0 ? width >> L : 1)$$

When **Corner Texel Mode** is enabled via the RENDER_SURFACE_STATE, the width of a 1D surface is calculated as shown below:

$$W_L = \text{MAX}(1, (W_{L-1} - 1) \gg 1) + 1$$

There is a restriction that the smallest map dimension is 2 texels for **Corner Texel Mode** ($W_0 > 1$)

Next, the aligned width parameter for each LOD "L" is computed. The "i" parameter is the horizontal alignment parameter set by a state field or defined as a constant, depending on the surface. The equation uses the i value that applies to the LOD being computed.

$$w_L = i * \text{ceil}\left(\frac{W_L}{i}\right)$$

Next, the offset to each LOD is determined. The offset has one dimension for 1D surfaces. The single element in the LOD_L vector is named $LODUL$.

$$\begin{aligned} LOD_0 &= (0) \\ LOD_1 &= (w_0) \\ LOD_2 &= (w_0 + w_1) \\ LOD_3 &= (w_0 + w_1 + w_2) \\ LOD_4 &= (w_0 + w_1 + w_2 + w_3) \\ &\dots \end{aligned}$$

Based on the above parameters and the U and R (pixel address and array index, respectively), and the bytes per pixel of the surface format (Bpp), the offset "u" in bytes from the base address of the surface is given by:

$$u = [(R * QPitch) + LODUL + U] * Bpp$$

Programming Note	
Context:	Packed YUV Surfaces
Packed YUV surface formats such as YCRCB_NORMAL, YCRCB_SWAPUVY etc. will be treated as 16bpp surface, not 32bpp, which may impact how they are laid out in memory.	

Tiling and Mip Tail for 1D Surfaces

There is no MIP Tail allowed for 1D surfaces because they are not allowed to be tiled. They must be declared as linear.

1D Alignment Requirements

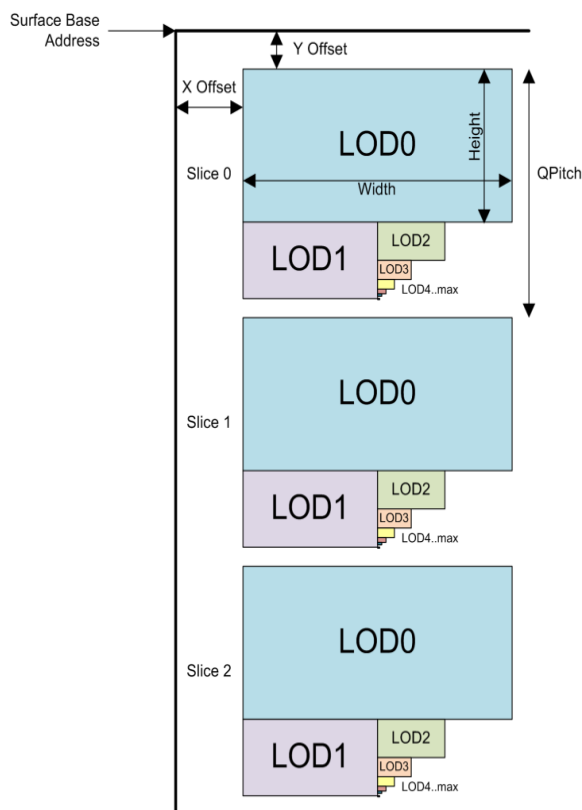
1D surfaces are not tiled, but laid out linearly in memory.

The **Horizontal Alignment** field in the Render Surface State is used to program alignment between MIPs.

2D Surfaces

2D surfaces represent two-dimensional bitmaps, which can also be mip-mapped and/or consist of array slices, effectively representing multiple 2D sub-surfaces within a single surface. The diagram below shows many of the parameters of a 2D surface or Arrayed 2D Surface and what they mean.

MIPs are NOT supported for MSAA surfaces.





All surface parameters are defined in terms of texels (agnostic to the surface format).

Surface Pitch defines the distance in bytes between rows of the surface and is a function of the **Width** of LOD. **QPitch** specifies the distance in rows between array slices and is a function of the **Height**. **QPitch** should allow at least enough space for any Mips that may be present.

There can also be non-zero offsets (**X_Offset** and **Y_Offset**) defined from the base address which can be used to provide padding or provide an offset to a lower-detail LOD.

There are limitations to the physical size of an LOD in the sampler texture cache. An LOD must be aligned to a cache-line except for some special cases related to Planar YUV surfaces. In general, the cache-alignment restriction implies there is a minimum height for an LOD of 4 texels. So, LODs which are smaller than 4 high are padded.

There are limitations to the physical size of an LOD in the texture cache. An LOD must be aligned to a cache-line except for some special cases related to Planar YUV surfaces. In general, the cache-alignment restriction implies there is a minimum height for an LOD of 4 texels. So, LODs which are smaller than 4 high are padded.

For tiled surfaces using Tile64 this restriction is not significant because the MIP tail will be used for smaller MIPs and the slots are a minimum of 256B. For Tile4, Mips smaller than 4 high start at the top of the region, and they are padded. This padding leads to a case where the smallest LOD starts "below" LOD1 vertically. Mip Tails can only be enabled for Tile64 surfaces.

Calculating Texel Location

This section describes how the texel location is calculated once the Surface State and LOD are known. A number of parameters are useful to determine where given pixels are located on the 2D surface. The width (**W_L**) and height (**H_L**) for each LOD "L" is computed by the formula:

$$W_L = ((\text{width} \gg L) > 0 ? \text{width} \gg L : 1)$$

$$H_L = ((\text{height} \gg L) > 0 ? \text{height} \gg L : 1)$$

When **Corner Texel Mode** is enabled via the `RENDER_SURFACE_STATE`, the width and height of a 2D surface are calculated as shown below:

$$W_L = \text{MAX}(1, (\text{W}_0 - 1) \gg L) + 1$$

$$H_L = \text{MAX}(1, (\text{H}_0 - 1) \gg L) + 1$$

There is a restriction that the smallest map dimension is 2 texels for **Corner Texel Mode** ($\text{W}_0 > 1$, $\text{H}_0 > 1$). This also applies to 2D arrays and 2D arrays viewed as cubes.

The LOD width and height for each subsequent LOD is one-half the previous LOD, with the minimum dimension being 1 texel. If the surface is multisampled and it is a depth or stencil surface or **Multisampled Surface Storage Format** in `SURFACE_STATE` is `MSFMT_DEPTH_STENCIL`, **W_L** and **H_L** must be adjusted as follows:

Number of Multisamples	$W_L =$	$H_L =$
2	$\text{ceiling}(W_L / 2) * 4$	H_L [no adjustment]
4	$\text{ceiling}(W_L / 2) * 4$	$\text{ceiling}(H_L / 2) * 4$
8	$\text{ceiling}(W_L / 2) * 8$	$\text{ceiling}(H_L / 2) * 4$
16	$\text{ceiling}(W_L / 2) * 8$	$\text{ceiling}(H_L / 2) * 8$

Next, aligned width, height, and depth parameters for each LOD "L" must be computed. The "i" and "j" parameters are horizontal and vertical alignment parameters set by state fields or defined as constants, depending on the surface. Depth has no alignment parameter (effectively it is 1).

The equation uses the i and j values that apply to the LOD being computed. The "p" and "q" parameters define the width and height in texels of the compression block for compressed surface formats. Both p and q are defined to equal 1 for uncompressed surface formats.

$$w_L = i * p * \text{ceil}\left(\frac{W_L}{i * p}\right)$$

$$h_L = j * q * \text{ceil}\left(\frac{H_L}{j * q}\right)$$

Once the height (h_i) and width (w_i) of each LOD is computed, the offset to each LOD can be determined. The offset is a vector with two dimensions. The elements in the LOD_L vector are named in order $LODU_L$, $LODV_L$.

LOD offset computation for when no Mip Tail is used or when $L < \text{Mip Tail Start LOD}$:

$$LOD_0 = (0,0)$$

$$LOD_1 = (0,h_0)$$

$$LOD_2 = (w_1,h_0)$$

$$LOD_3 = (w_1,h_0 + h_2)$$

$$LOD_4 = (w_1,h_0 + h_2 + h_3)$$

...

$$LOD_N = (w_1, h_0 + h_2 + h_3 \dots + h_{N-1})$$

Where $N = \text{MIP_COUNT}$ for the surface. As noted previous in this section, the value of $h_2 + h_3 \dots + h_{N-1}$ may be greater than h_1 due to alignment requirements.

Based on the above parameters and the U, V, and R (two-dimensional pixel address U/V and array index R), and the *bytes* per pixel of the surface format (Bpp), the offsets u in bytes and v in rows are given by:

$$u = (U + LODU_L) * \text{Bpp}$$

$$v = (R * \text{QPitch}) + LODV_L + V$$

For a description of how the Mip Tail is laid out and offsets into the Mip Tail are calculated see the sub-section on 2D Surface Layout for Mip Tails.



Programming Note

Context: Packed YUV Surfaces

Packed YUV surface formats such as YCRCB_NORMAL, YCRCB_SWAPUVY etc. will be treated as 16bpp surface, not 32bpp, which may impact how they are laid out in memory.

The two-dimensional offset into the surface (for non-MipTail cases) is defined by the u and v values computed above. The lower virtual address bits are determined by the following table, based on the bits of u and v. An *element* is defined as a pixel for uncompressed surface formats and a compression block for compressed surface formats. Empty bit positions indicate that the bit is not part of the tile swizzle and is filled in with equations given next (note that linear mode has all bits empty--there is no swizzling in linear mode).

The table below shows the mapping of u and v address bits within a tile for the supported tiling modes for a 2D surface. The u bits are a Byte address within a row and the v bits are a Row address within the tile.

Programming Note

Context: Tile64 64KB Address Swizzle

The tables below show the mapping of non-MSAA and MSAA. In the case of Arrayed Non-MSAA and MSAA the same mapping applies to each array slice.

See the definition of the **Surface QPitch** field in RENDER_SURFACE_STATE for how the height of the array slice is determined for MSAA and non-MSAA surfaces.

MSAA is only supported for Tile64.

Tile Mode	Bits per Element	TileID constants		Virtual Address Bits															
		Cv	Cu	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Tile64	128	6	10	v5	u9	u8	u7	v4	v3	u6	v2	u5	u4	v1	v0	u3	u2	u1	u0
	64	6	10	v5	u9	u8	u7	v4	v3	u6	v2	u5	u4	v1	v0	u3	u2	u1	u0
	32	7	9	v6	v5	u8	u7	v4	v3	u6	v2	u5	u4	v1	v0	u3	u2	u1	u0
	16	7	9	v6	v5	u8	u7	v4	v3	u6	v2	u5	u4	v1	v0	u3	u2	u1	u0
	8	8	8	v7	v6	v5	u7	v4	v3	u6	v2	u5	u4	v1	v0	u3	u2	u1	u0

Programming Note

Context:

Tile64 64KB Address Swizzle

Tile Mode	Bits per Element	TileID constants		Virtual Address Bits															
		Cv	Cu	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Tile64 2X MSAA	128	6	9	v5	u8	u7	u6	v4	v3	u5	v2	s0	u4	v1	v0	u3	u2	u1	u0
	64	6	9	v5	u8	u7	u6	v4	v3	u5	v2	s0	u4	v1	v0	u3	u2	u1	u0
	32	7	8	v6	v5	u7	u6	v4	v3	u5	v2	s0	u4	v1	v0	u3	u2	u1	u0
	16	7	8	v6	v5	u7	u6	v4	v3	u5	v2	s0	u4	v1	v0	u3	u2	u1	u0
	8	8	7	v7	v6	v5	u6	v4	v3	u5	v2	s0	u4	v1	v0	u3	u2	u1	u0

Tile Mode	Bits per Element	TileID constants		Virtual Address Bits															
		Cv	Cu	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Tile64 4X MSAA 8X MSAA* 16X MSAA* *requires 2 or 4 sample planes	128	5	9	v4	u8	u7	u6	v3	v2	u5	s1	s0	u4	v1	v0	u3	u2	u1	u0
	64	5	9	v4	u8	u7	u6	v3	v2	u5	s1	s0	u4	v1	v0	u3	u2	u1	u0
	32	6	8	v5	v4	u7	u6	v3	v2	u5	s1	s0	u4	v1	v0	u3	u2	u1	u0
	16	6	8	v5	v4	u7	u6	v3	v2	u5	s1	s0	u4	v1	v0	u3	u2	u1	u0
	8	7	7	v6	v5	v4	u6	v3	v2	u5	s1	s0	u4	v1	v0	u3	u2	u1	u0

Tile Mode	Bits per Element	TileID constants		Virtual Address Bits															
		Cv	Cu	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Tile4	all	5	7					v4	v3	u6	v2	u5	u4	v1	v0	u3	u2	u1	u0
Linear	all	0	0																
TileX	all	3	9					v2	v1	v0	u8	u7	u6	u5	u4	u3	u2	u1	u0



The TileID fills the upper bits of the virtual address (starting with the lowest blank bit in the above table):

$$\text{TileID} = (v \gg Cv) * (\text{Pitch} \gg Cu) + (u \gg Cu)$$

Where Pitch is the **Surface_Pitch** field from RENDER_SURFACE_STATE.

Note: Multisampled CMS and UMS surfaces use a modified address bit swizzling table rather than the one above. Refer to the *Multisampled2D Surfaces* section for details.

Tiling and Mip Tails for 2D Surfaces

When surface is Tiled using Tile64, a 2D surface can contain a Mip Tail for smaller Mip sizes.

When LOD (L) is less than the **Mip Tail Start LOD** (S) declared in the Surface State the offset to the start of LOD is calculated as shown above.

If the LOD is greater than or equal to **Mip Tail Start LOD** field in the surface state then the MIP Tail layout below is used.

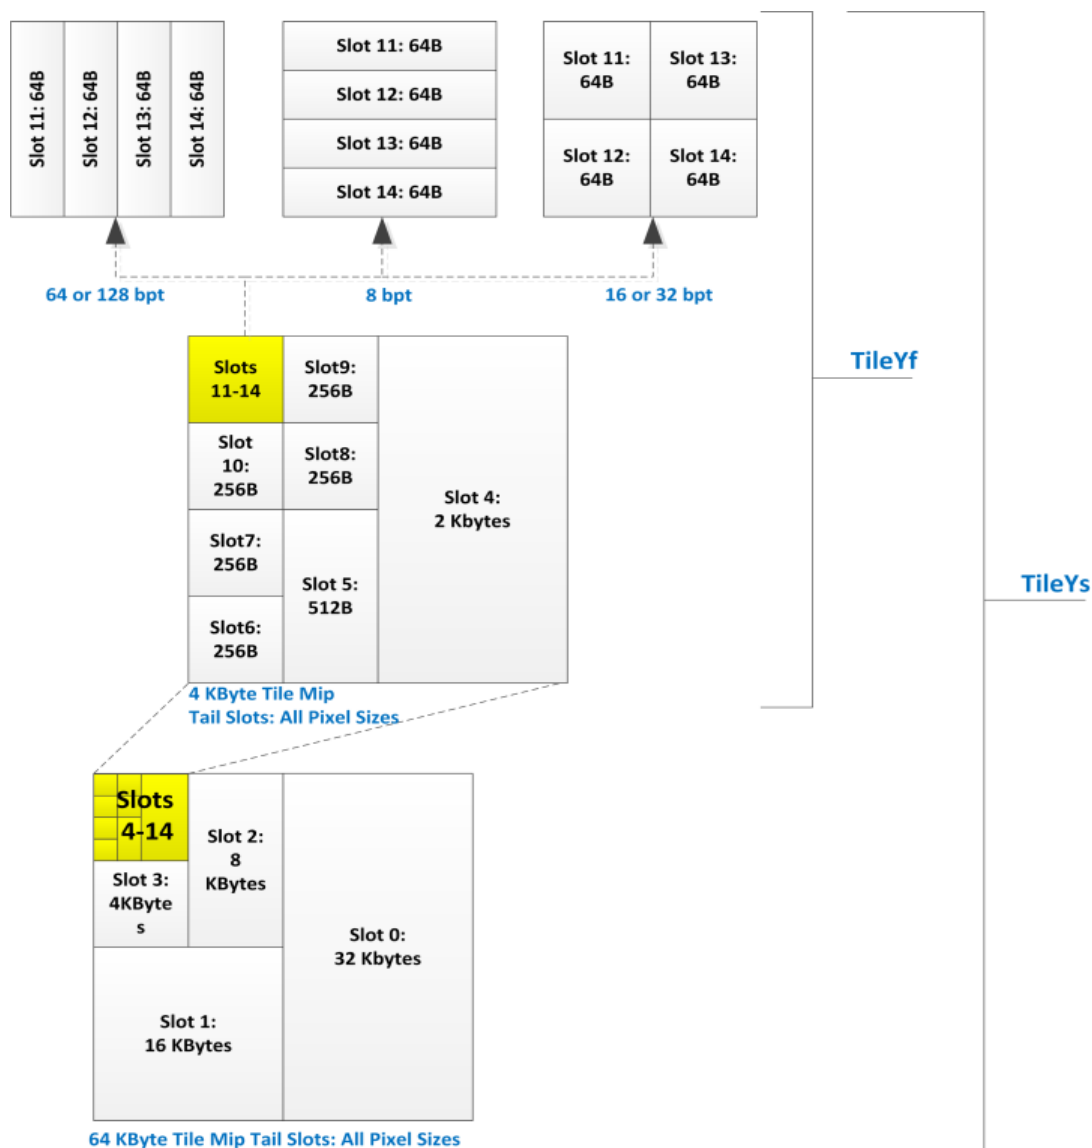
For tiled resources, the mip tail offset is given by the following, where s is the Mip Tail Start LOD:

$$\text{LOD}_s = (w_1, h_0 + h_2 + h_3 + \dots + h_{s-1})$$

The LOD's in the Mip Tail are arranged differently than the other LOD's.

The diagram below shows the 64KB TileYS Mip Tail layout of LODs within it, with "slots" indicating the LOD contained within (slot 0 corresponds to LOD_s above). LOD's are aligned to the upper left corner of the space available. The block marked "Slots 4-14" is a 4KB tile arrangement as shown. Within this 4KB tile slots 11 thru 14 are arranged differently depending on the number bits per texel (bpt).

A TileYf (4KByte) Mip Tail will start with the 4KByte tile shown, but the slots will be renumbered to start at Slot0 rather than Slot4. The layout of slots 11 through 14 remain the same. Note that Slots 12-14 are NOT 256-Byte aligned which is not compliant with the standard MIP Tail layout. These slots are not supported for Standard Tiling.



The offsets *into the Mip Tail tile* are given by the following table for each LOD in the Mip tail. Each entry in the table is a horizontal (M_U) and vertical (M_V) position (in texels) from the upper left corner of the Mip Tail. If $LOD \geq S$ (starting LOD for MIP Tail), then these Mip Tail offsets must be added to the LOD_U and LOD_V calculated above.

Note that many of the higher LODs are not possible given surface size constraints, but they are listed here for reference. The offsets given here need to be added to the LODs offset computed earlier to obtain the offset into the surface LOD_U .



Slot 11 is 256-byte aligned. Slots 12 through 14 are 64-byte aligned. MSAA is not supported for MIP Tails.

Tile64 LOD	128 bpe	64 bpe	32 bpe	16 bpe	8 bpe
s	(32,0)	(64,0)	(64,0)	(128,0)	(128,0)
s+1	(0,32)	(0,32)	(0,64)	(0,64)	(0,128)
s+2	(16,0)	(32,0)	(32,0)	(64,0)	(64,0)
s+3	(0,16)	(0,16)	(0,32)	(0,32)	(0,64)
s+4	(8,0)	(16,0)	(16,0)	(32,0)	(32,0)
s+5	(4, 8)	(8, 8)	(8, 16)	(16, 16)	(16, 32)
s+6	(0, 12)	(0, 12)	(0, 24)	(0, 24)	(0, 48)
s+7	(0, 8)	(0, 8)	(0, 16)	(0, 16)	(0, 32)
s+8	(4, 4)	(8, 4)	(8, 8)	(16, 8)	(16, 16)
s+9	(4, 0)	(8, 0)	(8, 0)	(16, 0)	(16, 0)
s+10	(0, 4)	(0, 4)	(0, 8)	(0, 8)	(0, 16)
s+11	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)
s+12	(1, 0)	(2, 0)	(0, 4)	(0, 4)	(0, 4)
s+13	(2, 0)	(4, 0)	(4, 0)	(8, 0)	(0, 8)
s+14	(3, 0)	(6, 0)	(4, 4)	(8, 4)	(0, 12)

The following table gives the maximum MIP size which can be included in the MIP Tail based on bpe. MIP Tail is only supported for Tile64.

Bits per Element	8	16	32	64	128
Maximum MIP Dimensions (in texels)	128x256	128x128	64x128	64x64	32x64

If the LOD is located in the MIP Tail then the equation for calculating the byte positions for u and v become:

$$u = (U + LODU_s + M_u) * Bpp$$

$$v = (R * QPitch) + LODV_s + M_v + V$$

where M_u and M_v are the offset parameters from the table above for the given slot in the MIP Tail.

Programming Note	
Context:	Lossless Compression and MIP Tail
Lossless compression must not be used on surfaces which have MIP Tail which contains MIPs for Slots greater than 11.	

Stencil Buffer Layout

This section details the layout of the stencil buffer.

1x Screen space(pixel)							
0	1	4	5	16	17	20	21
2	3	6	7	18	19	22	23
8	9	12	13	24	25	28	29
10	11	14	15	26	27	30	31
32	33	36	37	48	49	52	53
34	35	38	39	50	51	54	55
40	41	44	45	56	57	60	61
42	43	46	47	58	59	62	63

1x Physical space(pixel)															
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63

4x Screen space(pixel, sample)							
(0,0)	(1,0)	(0,1)	(1,1)	(4,0)	(5,0)	(4,1)	(5,1)
(2,0)	(3,0)	(2,1)	(3,1)	(6,0)	(7,0)	(6,1)	(7,1)
(0,2)	(1,2)	(0,3)	(1,3)	(4,2)	(5,2)	(4,3)	(5,3)
(2,2)	(3,2)	(2,3)	(3,3)	(6,2)	(7,2)	(6,3)	(7,3)
(8,0)	(9,0)	(8,1)	(9,1)	(12,0)	(13,0)	(12,1)	(13,1)
(10,0)	(11,0)	(10,1)	(11,1)	(14,0)	(15,0)	(14,1)	(15,1)
(8,2)	(9,2)	(8,3)	(9,3)	(12,2)	(13,2)	(12,3)	(13,3)
(10,2)	(11,2)	(10,3)	(11,3)	(14,2)	(15,2)	(14,3)	(15,3)

4x Physical space(pixel, sample)															
(0,0)	(1,0)	(2,0)	(3,0)	(0,1)	(1,1)	(2,1)	(3,1)	(0,2)	(1,2)	(2,2)	(3,2)	(0,3)	(1,3)	(2,3)	(3,3)
(4,0)	(5,0)	(6,0)	(7,0)	(4,1)	(5,1)	(6,1)	(7,1)	(4,2)	(5,2)	(6,2)	(7,2)	(4,3)	(5,3)	(6,3)	(7,3)
(8,0)	(9,0)	(10,0)	(11,0)	(8,1)	(9,1)	(10,1)	(11,1)	(8,2)	(9,2)	(10,2)	(11,2)	(8,3)	(9,3)	(10,3)	(11,3)
(12,0)	(13,0)	(14,0)	(15,0)	(12,1)	(13,1)	(14,1)	(15,1)	(12,2)	(13,2)	(14,2)	(15,2)	(12,3)	(13,3)	(14,3)	(15,3)



8x Screen space(pixel, sample)							
(0,0)	(1,0)	(0,1)	(1,1)	(0,4)	(1,4)	(0,5)	(1,5)
(2,0)	(3,0)	(2,1)	(3,1)	(2,4)	(3,4)	(2,5)	(3,5)
(0,2)	(1,2)	(0,3)	(1,3)	(0,6)	(1,6)	(0,7)	(1,7)
(2,2)	(3,2)	(2,3)	(3,3)	(2,6)	(3,6)	(2,7)	(3,7)
(4,0)	(5,0)	(4,1)	(5,1)	(4,4)	(5,4)	(4,5)	(5,5)
(6,0)	(7,0)	(6,1)	(7,1)	(6,4)	(7,4)	(6,5)	(7,5)
(4,2)	(5,2)	(4,3)	(5,3)	(4,6)	(5,6)	(4,7)	(5,7)
(6,2)	(7,2)	(6,3)	(7,3)	(6,6)	(7,6)	(6,7)	(7,7)

8x Physical space(pixel, sample)															
(0,0)	(1,0)	(2,0)	(3,0)	(0,1)	(1,1)	(2,1)	(3,1)	(0,2)	(1,2)	(2,2)	(3,2)	(0,3)	(1,3)	(2,3)	(3,3)
(0,4)	(1,4)	(2,4)	(3,4)	(0,5)	(1,5)	(2,5)	(3,5)	(0,6)	(1,6)	(2,6)	(3,6)	(0,7)	(1,7)	(2,7)	(3,7)
(4,0)	(5,0)	(6,0)	(7,0)	(4,1)	(5,1)	(6,1)	(7,1)	(4,4)	(5,4)	(6,4)	(7,4)	(4,5)	(5,5)	(6,5)	(7,5)
(4,2)	(5,2)	(6,2)	(7,2)	(4,3)	(5,3)	(6,3)	(7,3)	(4,6)	(5,6)	(6,6)	(7,6)	(4,7)	(5,7)	(6,7)	(7,7)

2D/CUBE Alignment Requirement

The vertical and horizontal alignment fields in the RENDER_SURFACE_STATE are ignored for Tile64. In the case of Tile64 the alignment requirements

are fixed and are provided for by the tables below for 2D and CUBE surface.

Tile Mode	Bits per Element	Horizontal Alignment	Vertical Alignment
Tile64	128	64	64
	64	128	64
	32	128	128
	16	256	128
	8	256	256

For MSFMT_MSS type multi-sampled Tile64 surfaces, the alignments given above must be divided by the appropriate value from the table below.

Number of Multisamples	Horizontal Alignment is divided by	Vertical Alignment is divided by
2	2	1
4	2	2
8	2	2
16	2	2

Multisampled 2D Surfaces

There are three types of multisampled surface layouts designated as follows:

- * **IMS** Interleaved Multisampled Surface
- * **CMS** Compressed Multisampled Surface
- * **UMS** Uncompressed Multisampled Surface

These surface layouts are described in the following sections.

Interleaved Multisampled Surfaces

These surfaces contain the samples in an interleaved fashion, with the underlying surface in memory having a height and width that is larger than the non-multisampled surface as follows:

- 4x MSAA: 2x width and 2x height of non-multisampled surface.
- 8x MSAA: 4x width and 2x height of non-multisampled surface.
- 16x MSAA: 4X width and 4X height of the non-multisampled surface.

When sampling from an IMS surface (e.g. Id2dms), the coordinates are automatically scaled to handle the increased physical size of the map.

The tables below shows the layout of 16-bit Depth (Z) values for different IMS formats. It shows layout of each 64-Byte chunk.



1X:

	bit 0							bit 127
Bytes 15:0	P(0,0)	P(1,0)	P(2,0)	P(3,0)	P(4,0)	P(5,0)	P(6,0)	P(7,0)
Bytes 16:31	P(0,1)	P(1,1)	P(2,1)	P(3,1)	P(4,1)	P(5,1)	P(6,1)	P(7,1)
Bytes 32:47	P(0,2)	P(1,2)	P(2,2)	P(3,2)	P(4,2)	P(5,2)	P(6,2)	P(7,2)
Bytes 48:63	P(0,3)	P(1,3)	P(2,3)	P(3,3)	P(4,3)	P(5,3)	P(6,3)	P(7,3)

2X:

	bit 0							bit 127
Bytes 15:0	P(0,0) s0	P(1,0) s0	P(0,0) s1	P(1,0) s1	P(2,0) s0	P(3,0) s0	P(2,0) s1	P(3,0) s1
Bytes 16:31	P(0,1) s0	P(1,1) s0	P(0,1) s1	P(1,1) s1	P(2,1) s0	P(3,1) s0	P(2,1) s1	P(3,1) s1
Bytes 32:47	P(0,2) s0	P(1,2) s0	P(0,2) s1	P(1,2) s1	P(2,2) s0	P(3,2) s0	P(2,2) s1	P(3,2) s1
Bytes 48:63	P(0,3) s0	P(1,3) s0	P(0,3) s1	P(1,3) s1	P(2,3) s0	P(3,3) s0	P(2,3) s1	P(3,3) s1

4X:

	bit 0							bit 127
Bytes 15:0	P(0,0) s0	P(1,0) s0	P(0,0) s1	P(1,0) s1	P(2,0) s0	P(3,0) s0	P(2,0) s1	P(3,0) s1
Bytes 16:31	P(0,1) s0	P(1,1) s0	P(0,1) s1	P(1,1) s1	P(2,1) s0	P(3,1) s0	P(2,1) s1	P(3,1) s1
Bytes 32:47	P(0,0) s2	P(1,0) s2	P(0,0) s3	P(1,0) s3	P(2,0) s2	P(3,0) s2	P(2,0) s3	P(3,0) s3
Bytes 48:63	P(0,1) s2	P(1,1) s2	P(0,1) s3	P(1,1) s3	P(2,1) s2	P(3,1) s2	P(2,1) s3	P(3,1) s3

8X:

	bit 0							bit 127
Bytes 15:0	P(0,0) s0	P(1,0) s0	P(0,0) s1	P(1,0) s1	P(0,0) s4	P(1,0) s4	P(0,0) s5	P(1,0) s5
Bytes 16:31	P(0,1) s0	P(1,1) s1	P(0,1) s1	P(1,1) s1	P(0,1) s4	P(1,1) s4	P(0,1) s5	P(1,1) s5
Bytes 32:47	P(0,0) s2	P(1,0) s2	P(0,0) s3	P(1,0) s3	P(0,0) s6	P(1,0) s6	P(0,0) s7	P(1,0) s7
Bytes 48:63	P(0,1) s2	P(1,1) s2	P(0,1) s3	P(1,1) s3	P(0,1) s6	P(1,1) s6	P(0,1) s7	P(1,1) s7

16X:

	bit 0							bit 127
Bytes 15:0	P(0,0) s0	P(1,0) s0	P(0,0) s1	P(1,0) s1	P(0,0) s4	P(1,0) s4	P(0,0) s5	P(1,0) s5
Bytes 16:31	P(0,1) s0	P(1,1) s0	P(0,1) s1	P(1,1) s1	P(0,1) s4	P(1,1) s4	P(0,1) s5	P(1,1) s5
Bytes 32:47	P(0,0) s2	P(1,0) s2	P(0,0) s3	P(1,0) s3	P(0,0) s6	P(1,0) s6	P(0,0) s7	P(1,0) s7
Bytes 48:63	P(0,1) s2	P(1,1) s2	P(0,1) s3	P(1,1) s3	P(0,1) s6	P(1,1) s6	P(0,1) s7	P(1,1) s7

	bit 0							bit 127
Bytes 64:79	P(0,0) s8	P(1,0) s8	P(0,0) s9	P(1,0) s9	P(0,0) s12	P(1,0) s12	P(0,0) s13	P(1,0) s13
Bytes 80:95	P(0,1) s8	P(1,1) s8	P(0,1) s9	P(1,1) s9	P(0,1) s12	P(1,1) s12	P(0,1) s13	P(1,1) s13
Bytes 96:111	P(0,0) s10	P(1,0) s10	P(0,0) s11	P(1,0) s11	P(0,0) s14	P(1,0) s14	P(0,0) s15	P(1,0) s15
Bytes 112:127	P(0,1) s10	P(1,1) s10	P(0,1) s11	P(1,1) s11	P(0,1) s14	P(1,1) s14	P(0,1) s15	P(1,1) s15



The table below shows the layout of 32-bit Depth (Z) values for different IMS formats. It shows layout of each 64-Byte chunk.

1X:

	bit 0			bit 127
Bytes 15:0	P(0,0)	P(1,0)	P(2,0)	P(3,0)
Bytes 16:31	P(0,1)	P(1,1)	P(2,1)	P(3,1)
Bytes 32:47	P(0,2)	P(1,2)	P(2,2)	P(3,2)
Bytes 48:63	P(0,3)	P(1,3)	P(2,3)	P(3,3)

2X:

	bit 0			bit 127
Bytes 15:0	P(0,0) s0	P(1,0) s0	P(0,0) s1	P(1,0) s1
Bytes 16:31	P(0,1) s0	P(1,1) s0	P(0,1) s1	P(1,1) s1
Bytes 32:47	P(0,2) s0	P(1,2) s0	P(0,2) s1	P(1,2) s1
Bytes 48:63	P(0,3) s0	P(1,3) s0	P(0,3) s1	P(1,3) s1

4X:

	bit 0			bit 127
Bytes 15:0	P(0,0) s0	P(1,0) s0	P(0,0) s1	P(1,0) s1
Bytes 16:31	P(0,1) s0	P(1,1) s0	P(0,1) s1	P(1,1) s1
Bytes 32:47	P(0,0) s2	P(1,0) s2	P(0,0) s3	P(1,0) s3
Bytes 48:63	P(0,1) s2	P(1,1) s2	P(0,1) s3	P(1,1) s3

8X:

	bit 0			bit 127
Bytes 15:0	P(0,0) s0	P(1,0) s0	P(0,0) s1	P(1,0) s1
Bytes 16:31	P(0,1) s0	P(1,1) s0	P(0,1) s1	P(1,1) s1
Bytes 32:47	P(0,0) s2	P(1,0) s2	P(0,0) s3	P(1,0) s3
Bytes 48:63	P(0,1) s2	P(1,1) s2	P(0,1) s3	P(1,1) s3

	bit 0			bit 127
Bytes 64:79	P(0,0) s4	P(1,0) s4	P(0,0) s5	P(1,0) s5
Bytes 80:95	P(0,1) s4	P(1,1) s4	P(0,1) s5	P(1,1) s5
Bytes 96:111	P(0,0) s6	P(1,0) s6	P(0,0) s7	P(1,0) s7
Bytes 112:127	P(0,1) s6	P(1,1) s6	P(0,1) s7	P(1,1) s7

16X:

	bit 0			bit 127
Bytes 15:0	P(0,0) s0	P(1,0) s0	P(0,0) s1	P(1,0) s1
Bytes 16:31	P(0,1) s0	P(1,1) s0	P(0,1) s1	P(1,1) s1
Bytes 32:47	P(0,0) s2	P(1,0) s2	P(0,0) s3	P(1,0) s3
Bytes 48:63	P(0,1) s2	P(1,1) s2	P(0,1) s3	P(1,1) s3

	bit 0			bit 127
Bytes 64:79	P(0,0) s4	P(1,0) s4	P(0,0) s5	P(1,0) s5
Bytes 80:95	P(0,1) s4	P(1,1) s4	P(0,1) s5	P(1,1) s5
Bytes 96:111	P(0,0) s6	P(1,0) s6	P(0,0) s7	P(1,0) s7
Bytes 112:127	P(0,1) s6	P(1,1) s6	P(0,1) s7	P(1,1) s7



	bit 0			bit 127
Bytes 128:143	P(0,0) s8	P(1,0) s8	P(0,0) s9	P(1,0) s9
Bytes 144:159	P(0,1) s8	P(1,1) s8	P(0,1) s9	P(1,1) s9
Bytes 160:175	P(0,0) s10	P(1,0) s10	P(0,0) s11	P(1,0) s11
Bytes 176:191	P(0,1) s10	P(1,1) s10	P(0,1) s11	P(1,1) s11

	bit 0			bit 127
Bytes 192:207	P(0,0) s12	P(1,0) s12	P(0,0) s13	P(1,0) s13
Bytes 208:223	P(0,1) s12	P(1,1) s12	P(0,1) s13	P(1,1) s13
Bytes 224:239	P(0,0) s14	P(1,0) s14	P(0,0) s15	P(1,0) s15
Bytes 240:255	P(0,1) s14	P(1,1) s14	P(0,1) s15	P(1,1) s15

The table below shows the layout of Depth Stencil values for different IMS formats. It shows layout of each 64-Byte chunk.

1X:

Bytes 0-7	P(0,0)	P(1,0)	P(0,1)	P(1,1)	P(2,0)	P(3,0)	P(2,1)	P(3,1)
Bytes 8-15	P(0,2)	P(1,2)	P(0,3)	P(1,3)	P(2,2)	P(3,2)	P(2,3)	P(3,3)
Bytes 16-23	P(4,0)	P(5,0)	P(4,1)	P(5,1)	P(6,0)	P(7,0)	P(6,1)	P(7,1)
Bytes 24-31	P(4,2)	P(5,2)	P(4,3)	P(5,3)	P(6,2)	P(7,2)	P(6,3)	P(7,3)
Bytes 32-39	P(0,4)	P(1,4)	P(0,5)	P(1,5)	P(2,4)	P(3,4)	P(2,5)	P(3,5)
Bytes 40-47	P(0,6)	P(1,6)	P(0,7)	P(1,7)	P(2,6)	P(3,6)	P(2,7)	P(3,7)
Bytes 48-55	P(4,4)	P(5,4)	P(4,5)	P(5,5)	P(6,4)	P(7,4)	P(6,5)	P(7,5)
Bytes 56-63	P(4,6)	P(5,6)	P(4,7)	P(5,7)	P(6,6)	P(7,6)	P(6,7)	P(7,7)

2X:

Bytes 0-7	P(0,0) s0	P(1,0) s0	P(0,1) s0	P(1,1) s0	P(0,0) s1	P(1,0) s1	P(0,1) s1	P(1,1) s1
Bytes 8-15	P(0,2) s0	P(1,2) s0	P(0,3) s0	P(1,3) s0	P(0,2) s1	P(1,2) s1	P(0,3) s1	P(1,3) s1
Bytes 16-23	P(2,0) s0	P(3,0) s0	P(2,1) s0	P(3,1) s0	P(2,0) s1	P(3,0) s1	P(2,1) s1	P(3,1) s1
Bytes 24-31	P(2,2) s0	P(3,2) s0	P(2,3) s0	P(3,3) s0	P(2,2) s1	P(3,2) s1	P(2,3) s1	P(3,3) s1
Bytes 32-39	P(0,4) s0	P(1,4) s0	P(0,5) s0	P(1,5) s0	P(0,4) s1	P(1,4) s1	P(0,5) s1	P(1,5) s1
Bytes 40-47	P(0,6) s0	P(1,6) s0	P(0,7) s0	P(1,7) s0	P(0,6) s1	P(1,6) s1	P(0,7) s1	P(1,7) s1
Bytes 48-55	P(2,4) s0	P(3,4) s0	P(2,5) s0	P(3,5) s0	P(2,4) s1	P(3,4) s1	P(2,5) s1	P(3,5) s1
Bytes 56-63	P(2,6) s0	P(3,6) s0	P(2,7) s0	P(3,7) s0	P(2,6) s1	P(3,6) s1	P(2,7) s1	P(3,7) s1

4X:

Bytes 0-7	P(0,0) s0	P(1,0) s0	P(0,1) s0	P(1,1) s0	P(0,0) s1	P(1,0) s1	P(0,1) s1	P(1,1) s1
Bytes 8-15	P(0,0) s2	P(1,0) s2	P(0,1) s2	P(1,1) s2	P(0,0) s3	P(1,0) s3	P(0,1) s3	P(1,1) s3
Bytes 16-23	P(2,0) s0	P(3,0) s0	P(2,1) s0	P(3,1) s0	P(2,0) s1	P(3,0) s1	P(2,1) s1	P(3,1) s1
Bytes 24-31	P(2,0) s2	P(3,0) s2	P(2,1) s2	P(3,1) s2	P(2,0) s3	P(3,0) s3	P(2,1) s3	P(3,1) s3
Bytes 32-39	P(0,2) s0	P(1,2) s0	P(0,3) s0	P(1,3) s0	P(0,2) s1	P(1,2) s1	P(0,3) s1	P(1,3) s1
Bytes 40-47	P(0,2) s2	P(1,2) s2	P(0,3) s2	P(1,3) s2	P(0,2) s3	P(1,2) s3	P(0,3) s3	P(1,3) s3
Bytes 48-55	P(2,2) s0	P(3,2) s0	P(2,3) s0	P(3,3) s0	P(2,2) s1	P(3,2) s1	P(2,3) s1	P(3,3) s1
Bytes 56-63	P(2,2) s2	P(3,2) s2	P(2,3) s2	P(3,3) s2	P(2,2) s3	P(3,2) s3	P(2,3) s3	P(3,3) s3



8X:

Bytes 0-7	P(0,0) s0	P(1,0) s0	P(0,1) s0	P(1,1) s0	P(0,0) s1	P(1,0) s1	P(0,1) s1	P(1,1) s1
Bytes 8-15	P(0,0) s2	P(1,0) s2	P(0,1) s2	P(1,1) s2	P(0,0) s3	P(1,0) s3	P(0,1) s3	P(1,1) s3
Bytes 16-23	P(0,0) s4	P(1,0) s4	P(0,1) s4	P(1,1) s4	P(0,0) s5	P(1,0) s5	P(0,1) s5	P(1,1) s5
Bytes 24-31	P(0,0) s6	P(1,0) s6	P(0,1) s6	P(1,1) s6	P(0,0) s7	P(1,0) s7	P(0,1) s7	P(1,1) s7
Bytes 32-39	P(0,2) s0	P(1,2) s0	P(0,3) s0	P(1,3) s0	P(0,2) s1	P(1,2) s1	P(0,3) s1	P(1,3) s1
Bytes 40-47	P(0,2) s2	P(1,2) s2	P(0,3) s2	P(1,3) s2	P(0,2) s3	P(1,2) s3	P(0,3) s3	P(1,3) s3
Bytes 48-55	P(0,2) s4	P(1,2) s4	P(0,3) s4	P(1,3) s4	P(0,2) s5	P(1,2) s5	P(0,3) s5	P(1,3) s5
Bytes 56-63	P(0,2) s6	P(1,2) s6	P(0,3) s6	P(1,3) s6	P(0,2) s7	P(1,2) s7	P(0,3) s7	P(1,3) s7

16X:

Bytes 0-7	P(0,0) s0	P(1,0) s0	P(0,1) s0	P(1,1) s0	P(0,0) s1	P(1,0) s1	P(0,1) s1	P(1,1) s1
Bytes 8-15	P(0,0) s2	P(1,0) s2	P(0,1) s2	P(1,1) s2	P(0,0) s3	P(1,0) s3	P(0,1) s3	P(1,1) s3
Bytes 16-23	P(0,0) s4	P(1,0) s4	P(0,1) s4	P(1,1) s4	P(0,0) s5	P(1,0) s5	P(0,1) s5	P(1,1) s5
Bytes 24-31	P(0,0) s6	P(1,0) s6	P(0,1) s6	P(1,1) s6	P(0,0) s7	P(1,0) s7	P(0,1) s7	P(1,1) s7
Bytes 32-39	P(0,0) s8	P(1,0) s8	P(0,1) s8	P(1,1) s8	P(0,0) s9	P(1,0) s9	P(0,1) s9	P(1,1) s9
Bytes 40-47	P(0,0) s10	P(1,0) s10	P(0,1) s10	P(1,1) s10	P(0,0) s11	P(1,0) s11	P(0,1) s11	P(1,1) s11
Bytes 48-55	P(0,0) s12	P(1,0) s12	P(0,1) s12	P(1,1) s12	P(0,0) s13	P(1,0) s13	P(0,1) s13	P(1,1) s13
Bytes 56-63	P(0,0) s14	P(1,0) s14	P(0,1) s14	P(1,1) s14	P(0,0) s15	P(1,0) s15	P(0,1) s15	P(1,1) s15

Compressed Multisampled Surfaces

Multisampled render targets can be compressed. If **Auxiliary Surface Mode** in SURFACE_STATE is set to AUX_CCS, hardware handles the compression using a software-invisible algorithm. However, performance optimizations in the multisample resolve kernel using the sampling engine are possible if the internal format of these surfaces is understood by software. This section documents the formats of the Multisample Control Surface (MCS) and Multisample Surface (MSS).

MCS Surface

The MCS surface consists of one element per pixel, with the element size being an 8-bit unsigned integer value for 4x multisampled surfaces, a 32-bit unsigned integer value for 8x multisampled surfaces and a 64-bit unsigned integer value for 16x multisampled surface. Each field within the element indicates which sample slice (SS) the sample resides on.

2x MCS

The 2x MCS is 8 bits per pixel. The 8 bits are encoded as follows:

7:2	1	0
reserved	sample 1 SS	sample 0 SS

Each 1-bit field indicates which sample slice (SS) the sample's color value is stored. An MCS value of 0x00 indicates that both samples are stored in sample slice 0 (thus have the same color). This is the fully compressed case. An MCS value of 0x03 indicates that all samples in the pixel are in the clear state, and none of the sample slices are valid. The pixel's color must be replaced with the surface's clear value.

4x MCS

The 4x MCS is 8 bits per pixel. The 8 bits are encoded as follows:

7:6	5:4	3:2	1:0
sample 3 SS	sample 2 SS	sample 1 SS	sample 0 SS

Each 2-bit field indicates which sample slice (SS) the sample's color value is stored. An MCS value of 0x00 indicates that all four samples are stored in sample slice 0 (thus all have the same color). This is the fully compressed case. An MCS value of 0xff indicates that all samples in the pixel are in the clear state, and none of the sample slices are valid. The pixel's color must be replaced with the surface's clear value. See the section below on Clear Pixel Conditions for additional encoding information.

8x MCS

Extending the mechanism used for the 4x MCS to 8x requires 3 bits per sample times 8 samples, or 24 bits per pixel. The 24-bit MCS value per pixel is placed in a 32-bit footprint, with the upper 8 bits unused as shown below. See the section below on Clear Pixel Conditions for additional encoding information.

31:24	23:21	20:18	17:15	14:12	11:9	8:6	5:3	2:0
reserved (MBZ)	sample 7 SS	sample 6 SS	sample 5 SS	sample 4 SS	sample 3 SS	sample 2 SS	sample 1 SS	sample 0 SS

16x MCS

The 16x MCS is 64 bits per pixel. The 64 bits are encoded as follows:

63:60	59:56	55:52	51:48	47:44	43:40	39:36	35:32
sample 15 SS	sample 14 SS	sample 13 SS	sample 12 SS	sample 11 SS	sample 10 SS	sample 9 SS	sample 8 SS



31:28	27:24	23:20	19:16	15:12	11:8	7:4	3:0
sample 7 SS	sample 6 SS	sample 5 SS	sample 4 SS	sample 3 SS	sample 2 SS	sample 1 SS	sample 0 SS

Other than this, the 16x algorithm is the same as the 8x algorithm. The MCS value indicating clear state is 0xffffffff_ffffffff. See the section below on Clear Pixel Conditions for additional encoding information.

Clear Pixel Conditions

The MCS format allows for the encoding of clear value for one or more planes of the multi-sampled surface. A value of all 1's for defined MCS bits indicates that all planes of the multi-sampled surface are clear. For example, a value of 0x3 for 2X MSAA MCS byte means that both planes of the pixel are clear. Likewise, a value of 0xff for X4, 0xffffffff for X8 and 0xffffffff_ffffffff for X16 MSAA means that all planes of the pixel are clear.

In the case where not all planes are clear, but at least 2 planes are clear the encoding of the MCS given above is changed. If the MCS value for plane 0 is non-zero, then all planes which are at all 1's are clear and all other planes are referencing the plane indicated by their respective MCS value minus 1. For example, a 4X MSAA MCS value of 01 10 11 11 means that MCS 0 is referencing plane 0, and MCS 1 is referencing plane 1, and MCS 2 and 3 are clear.

MSS Surface

The physical MSS surface is stored identically to a 2D array surface, with the height and width matching the pixel dimensions of the logical multisampled surface. The number of array slices in the physical surface is 2, 4, 8, or 16 times that of the logical surface (depending on the number of multisamples). Sample slices belonging to the same logical surface array slice are stored in adjacent physical slices. The sampling engine ld2dss message gives direct access to a specific sample slice.

Tiling for CMS and UMS Surfaces

Multisampled CMS and UMS use a modified table from non-multisampled 2D surfaces.

TileY,

TileX, TileW, and Linear: Treat as 2D array, with the array index "R" modified as follows. "n" is the number of multisamples, "ss" is the sample slice index with range 0..n-1.

$$R_{(new)} = (R_{(old)} \ll \log_2(n)) | ss$$

See 2D Surfaces main section for the layout of MSAA tiles. MSAA is only supported with Tile64.

TileYF: In addition to u and v, the sample slice index "ss" is included in the address swizzling according to the following table. Because of this, the mip tail holds one less LOD for each successive number of multisamples. Refer to the mip tail table in the previous section for behavior of the mip tail for each number of multisamples.

Uncompressed Multisampled Surfaces

UMS surfaces similar to CMS, except that the **Auxiliary Surface Mode** is set to AUX_NONE, meaning that there is no MCS surface. UMS contains only an MSS surface, where each sample is stored on its sample slice (SS) of the same index.

Programming Note	
Context:	3D Sampler
See 3D Sampler Messages section for a description of how the Id2dms and Id2dms_w messages work for UMS surfaces	

Cube Surfaces

The 3D Pipeline supports *cubic environment maps*, conceptually arranged as a cube surrounding the origin of a 3D coordinate system aligned to the cube faces. These maps can be used to supply texel (color/alpha) data of the environment in any direction from the enclosed origin, where the direction is supplied as a 3D "vector" texture coordinate. These cube maps can also be mipmapped.


Each texture map level is represented as a group of six, square *cube face* texture surfaces. The faces are identified by their relationship to the 3D texture coordinate system. The subsections below describe the cube maps as described at the API as well as the memory layout dictated by the hardware.

The diagram below describes the cube map faces as they are defined at the DirectX API. It shows the axes on the faces as they would be seen from the inside (at the origin).

The 3D sampler converts the incoming U,V,R coordinates on the sampler.

This will be looking directly at face 4, the +z -face. Y is up by default.

DirectX Cube Map Definition



B.6687-01



The coordinates on each face are relative to the center of the cube, and they range from -1.0 to 1.0 rather than the normal 0 to 1.0 normalized coordinate system in a 2D array surface.

Each face has a corresponding face identifier "f" as indicated in the following table:

face	face identifier "f"
+x	0
-x	1
+y	2
-y	3
+z	4
-z	5

A cube surface is stored in memory the same as a 2D array, with the face identifier "f" and array index "ai" being transformed into the "R" coordinate used in storing 2D arrays using the following equation:

$$R = (ai * 6) + f$$


Refer to the "2D Surfaces" section for details on how 2D arrays are stored.

3D Surfaces

Multiple texture map surfaces (and their respective mipmap chains) can be arranged into a structure known as a Texture3D (volume) texture. A volume texture map consists of many *planes* of 2D texture maps.

See *Sampler* for a description of how volume textures are used.

Volume Texture Map



B6688-01

Surface Pitch defines the distance in bytes between rows of the surface. **Surface QPitch** specifies the distance in rows between R-slices. QPitch should allow at least enough space for any mipmaps that may be present.

A number of parameters are useful to determine where given pixels are located on the 3D surface. First, the width, height, and depth for each LOD "L" is computed:

$$W_L = ((width \gg L) > 0 ? width \gg L : 1)$$

$$H_L = ((height \gg L) > 0 ? height \gg L : 1)$$

$$D_L = ((depth \gg L) > 0 ? depth \gg L : 1)$$

When **Corner Texel Mode** is enabled via the RENDER_SURFACE_STATE, the width and height of a 3D surface are calculated as shown below:

$$W_L = \text{MAX}(1, (W_0 - 1) \gg L) + 1$$

$$H_L = \text{MAX}(1, (H_0 - 1) \gg L) + 1$$

$$D_L = \text{MAX}(1, (D_0 - 1) \gg L) + 1$$

There is a restriction that the smallest map dimension is 2 texels for **Corner Texel Mode** ($W_0 > 1$, $H_0 > 1$, $D_0 > 1$)

Next, aligned width, height, and depth parameters for each LOD "L" are computed. The "i", "j", and "k" parameters are the horizontal, vertical, and depth alignment parameters set by state fields or defined as constants. The alignment parameters may change at one point in the mip chain based on **Mip Tail Start LOD**. The equation uses the i/j values that apply to the LOD being computed. The "p", "q", and "s" parameters define the width, height, and depth in texels of the compression block for compressed surface formats. These are all defined to equal 1 for uncompressed surface formats.

$$w_L = i * \text{ceil}\left(\frac{W_L}{i}\right)$$

$$h_L = j * \text{ceil}\left(\frac{H_L}{j}\right)$$

$$d_L = D_L$$

$$d'_L = k * s * \text{ceil}\left(\frac{D_L}{k * s}\right)$$

Next, the offset to each LOD is determined. The offset is a vector with three dimensions. The elements in the LOD_L vector are named in order LOD_U , LOD_V , LOD_R .

LOD offset computation for **Tiled Resource Mode** == TR_NONE or when $L < \text{Mip Tail Start LOD}$:

$$\text{LOD}_0 = (0, 0, 0)$$

$$\text{LOD}_1 = (0, h_0, 0)$$

$$\text{LOD}_2 = (w_1, h_0, 0)$$

$$\text{LOD}_3 = (w_1, h_0 + h_2, 0)$$

$$\text{LOD}_4 = (w_1, h_0 + h_2 + h_3, 0)$$

...

For the Primary Surface

Based on the above parameters and the U, V, and R (three-dimensional pixel address), and the bytes per pixel of the surface format (Bpp), the offsets u in bytes, v in rows, and r in slices are given by:

$$u = [U + \text{LOD}_U] * \text{Bpp}$$

$$v = \text{LOD}_V + V$$

$$r = \text{LOD}_R + R$$

Programming Note	
Context:	Packed YUV Surfaces
Packed YUV surface formats such as YCRCB_NORMAL, YCRCB_SWAPUVY etc. will be treated as 16bpp surface, not 32bpp, which may impact how they are laid out in memory.	

The three-dimensional offset into the surface is defined by the u, v, and r values computed above. The lower virtual address bits are determined by the following table, based on the bits of u, v, and r. An *element* is defined as a pixel for uncompressed surface formats and a compression block for compressed surface formats.

Empty bit positions indicate that the bit is not part of the tile swizzle and is filled in with the equations given next (note that linear mode has all bits empty--there is no swizzling in linear mode).

The table below is enabled by use of the Tile Address Mapping Mode bi in **RENDER_SURFACE_STATE**.

This mapping must never be used.

Tile Mode	Bits per Element	TileID constants			Virtual Address Bits															
		Cr	Cv	Cu	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Tile64	128	4	4	8	r3	r2	v3	u7	u6	u5	r1	v2	r0	u4	v1	v0	u3	u2	u1	u0
	64	4	4	8	r3	r2	v3	u7	u6	u5	r1	v2	r0	u4	v1	v0	u3	u2	u1	u0
	32	4	5	7	r3	r2	v4	u6	v3	u5	r1	v2	r0	u4	v1	v0	u3	u2	u1	u0
	16	5	5	6	r4	r3	r2	v4	v3	u5	r1	v2	r0	u4	v1	v0	u3	u2	u1	u0
	8	5	5	6	r4	r3	r2	u5	v4	v3	r1	v2	r0	u4	v1	v0	u3	u2	u1	u0
Linear	all	0	0	0																

Tile Mode	Bits per Element	TileID constants			Virtual Address Bits															
		Cr	Cv	Cu	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Tile4	all	N/A	5	7					v4	v3	u6	v2	u5	u4	v1	v0	u3	u2	u1	u0
Linear	all	N/A	0	0																

Note that Tile4 does not combine different Q slices (for 3D or MSAA) in a single 4K tile. Each tile contains a single Q or Sub-pixel.

The TileID fills the upper bits of the virtual address (starting with the lowest blank bit in the above table):

$$\text{TileID} = [(r \gg Cr) * (QPitch \gg Cv) + (v \gg Cv)] * (Pitch \gg Cu) + (u \gg Cu)$$

Tiling and Mip Tails for 3D Surfaces

For tiled surfaces where Tiled Resource Mode != TR_NONE, the surface may contain a mip tail. The Mip tail offset is given by the following, where S is the Mip Tail Start LOD:

$$LOD_S = (w_1, h_0 + h_2 + h_3 + \dots + h_{S-1}, 0)$$

The offsets into the mip tail tile are given by the following table for each LOD in the mip tail. Note that many of the higher LODs are not possible given surface size constraints, but they are listed here for reference. The offsets given here need to be added to the LODs offset computed earlier to obtain the offset into the surface LOD_L .

Slot Number	Slot Size In Bytes	128-bit 3D Offset	64-bit 3D Offset	32-bit 3D Offset	16-bit 3D Offset	8-bit 3D Offset
0	32KB	(8, 0, 0)	(16, 0, 0)	(16, 0, 0)	(16, 0, 0)	(32, 0, 0)
1	16KB	(0, 8, 0)	(0, 8, 0)	(0, 16, 0)	(0, 16, 0)	(0, 16, 0)
2	8KB	(0, 0, 8)	(0, 0, 8)	(0, 0, 8)	(0, 0, 16)	(0, 0, 16)
3	4KB	(4, 0, 0)	(8, 0, 0)	(8, 0, 0)	(8, 0, 0)	(16, 0, 0)
4	2KB	(0, 4, 0)	(0, 4, 0)	(0, 8, 0)	(0, 8, 0)	(0, 8, 0)
5	512B	(2, 0, 4)	(4, 0, 4)	(4, 0, 4)	(0, 4, 8)	(0, 4, 8)
6	256B	(1, 0, 4)	(2, 0, 4)	(0, 4, 4)	(0, 0, 12)	(0, 0, 12)
7	256B	(0, 0, 4)	(0, 0, 4)	(0, 0, 4)	(0, 0, 8)	(0, 0, 8)
8	256B	(3, 0, 0)	(6, 0, 0)	(4, 4, 0)	(0, 4, 4)	(0, 4, 4)
9	256B	(2, 0, 0)	(4, 0, 0)	(4, 0, 0)	(0, 4, 0)	(0, 4, 0)
10	256B	(1, 0, 0)	(2, 0, 0)	(0, 4, 0)	(0, 0, 4)	(0, 0, 4)
11	64B	(0, 0, 0)	(0, 0, 0)	(0, 0, 0)	(0, 0, 0)	(0, 0, 0)
[12]	64B	(0, 0, 1)	(0, 0, 1)	(0, 0, 1)	(0, 0, 1)	(0, 0, 1)
[13]	64B	(0, 0, 2)	(0, 0, 2)	(0, 0, 2)	(0, 0, 2)	(0, 0, 2)
[14]	64B	(0, 0, 3)	(0, 0, 3)	(0, 0, 3)	(0, 0, 3)	(0, 0, 3)



The following table gives the maximum MIP size which can be included in the MIP Tail based on bpe. MIP Tail is only supported for Tile64.

Bits per Element	8	16	32	64	128
Maximum MIP Dimensions (in texels)	32x32x32	16x32x32	16x32x16	16x16x16	8x16x16

3D Alignment Requirements

The vertical and horizontal alignment fields in the RENDER_SURFACE_STATE are ignored for Tile64.

In the case of Tile64 the alignment requirements are fixed and are provided for by the tables below for 3D (volumetric) surfaces.

Tile Mode	Bits per Element	Horizontal Alignment	Vertical Alignment	Depth Alignment
Tile64	128	16	16	16
	64	32	16	16
	32	32	32	16
	16	32	32	32
	8	64	32	32

Surface Padding Requirements

This section covers the requirements for padding around surfaces stored in memory, as there are cases where the device will overfetch beyond the bounds of the surface due to implementation of caches and other hardware structures.

Sampling Engine Surfaces

The sampling engine accesses texels outside of the surface if they are contained in the same cache line as texels that are within the surface. These texels will not participate in any calculation performed by the sampling engine and will not affect the result of any sampling engine operation, however if these texels lie outside of defined pages in the GTT, a GTT error will result when the cache line is accessed. In order to avoid these GTT errors, "padding" at the bottom and right side of a sampling engine surface is sometimes necessary.

It is possible that a cache line will straddle a page boundary if the base address or pitch is not aligned. All pages included in the cache lines that are part of the surface must map to valid GTT entries to avoid errors. To determine the necessary padding on the bottom and right side of the surface, refer to the table in Alignment Unit Size section for the i and j parameters for the surface format in use. The surface must then be extended to the next multiple of the alignment unit size in each dimension, and all texels contained in this extended surface must have valid GTT entries.

For example, suppose the surface size is 15 texels by 10 texels and the alignment parameters are $i=4$ and $j=2$. In this case, the extended surface would be 16 by 10. Note that these calculations are done in texels, and must be converted to bytes based on the surface format being used to determine whether additional pages need to be defined.

For compressed textures (BC*, FXT1, ETC*, and EAC* surface formats), padding at the bottom of the surface is to an even compressed row. This is equivalent to a multiple of $2q$, where q is the compression block height in texels. Thus, for padding purposes, these surfaces behave as if $j = 2q$ only for surface padding purposes. The value of j is still equal to q for mip level alignment and QPitch calculation. For cube surfaces, an additional two rows of padding are required at the bottom of the surface. This must be ensured regardless of whether the surface is stored tiled or linear. This is due to the potential rotation of cache line orientation from memory to cache.

For packed YUV, 96 bpt, 48 bpt, and 24 bpt surface formats, additional padding is required. These surfaces require an extra row plus 16 bytes of padding at the bottom in addition to the general padding requirements.

For linear surfaces, additional padding of 64 bytes is required at the bottom of the surface. This is in addition to the padding required above.

Programming Note	
Context:	Sampling Engine Surfaces.
For SURFTYPE_BUFFER, SURFTYPE_1D, and SURFTYPE_2D non-array, non-MSAA, non-mip-mapped surfaces in linear memory, the only padding requirement is to the next aligned 64-byte boundary beyond the end of the surface. The rest of the padding requirements documented above do not apply to these surfaces.	

Programming Note	
Context:	Sampling Engine Surfaces
For all surface types other than non-mipmapped non-arrayed 2D, 1D, and Buffer, when using linear mode and surface Height %4 != 0, the surface must be padded with $4 - (\text{Height} \% 4) * \text{Surface_Pitch}$ bytes to avoid fetching outside of allocated memory.	

Render Target and Media Surfaces

The data port accesses data (pixels) outside of the surface if they are contained in the same cache request as pixels that are within the surface. These pixels will not be returned by the requesting message, however if these pixels lie outside of defined pages in the GTT, a GTT error will result when the cache request is processed. In order to avoid these GTT errors, "padding" at the bottom of the surface is sometimes necessary.



Address Tiling Function Introduction

When dealing with memory operands (e.g., graphics surfaces) that are inherently rectangular in nature, certain functions within the graphics device support the storage/access of the operands using alternative (tiled) memory formats to increase performance. This section describes these memory storage formats, why and when they should be used, and the behavioral mechanisms within the device to support them.

Legacy Tiling Modes:

- **TileX** : Used primarily for display surfaces.

Additional Tiling Modes

- **Tile4**: 4KB tiling mode based on previously-supported TileY
- **Tile64**: 64KB tiling mode which support standard-tiling including Mip Tails

The Tile64 tile is made up of a 4x4 matrix of 4KB tiles. The Tile64 tile, in general, has a different height and width (in pixels) determined by the bits per element (bpe) of the surface format in attempt to keep the tile as close to square as possible.

The Tile4 tile has a fixed height of 32 pixels for all bpe (8,16,32,64,128), but the width in pixels varies accordingly.


All tile modes are enabled via the **Tile Mode** field in the `RENDER_SURFACE_STATE`.

Tiled surface base addresses must be tile aligned (64KB aligned for Tile64, 4KB aligned for all other tile modes). An exception to this tile alignment is when a `SURFACE_STATE` describes a single MIP within the MIP Tail of another surface, using a 64-bit or 128-bit **Surface Format**--then **Surface Base Address** can refer directly to the given MIP (e.g. to write to a non-renderable **Surface Format** by re-describing as an alternative surface).

Linear vs Tiled Storage

Regardless of the memory storage format, "rectangular" memory operands have a specific *width* and *height*, and are considered as residing within an enclosing rectangular region whose width is considered the *pitch* of the region and surfaces contained within. Surfaces stored within an enclosing region must have widths less than or equal to the region pitch (indeed the enclosing region may coincide exactly with the surface). *Rectangular Memory Operand Parameters* shows these parameters.


Rectangular Memory Operand Parameters



B.6690-01

The simplest storage format is the *linear* format (see *Linear Surface Layout*), where each row of the operand is stored in sequentially increasing memory locations. If the surface width is less than the enclosing region's pitch, there will be additional memory storage between rows to accommodate the region's pitch. The pitch of the enclosing region determines the distance (in the memory address space) between vertically-adjacent operand elements (e.g., pixels, texels).

Linear Surface Layout



B.6691-01

The linear format is best suited for 1-dimensional row-sequential access patterns (e.g., a display surface where each scanline is read sequentially). Here the fact that one object element may reside in a different memory page than its vertically-adjacent neighbors is not significant; all that matters is that horizontally-adjacent elements are stored contiguously. However, when a device function needs to access a 2D subregion within an operand (e.g., a read or write of a 4x4 pixel span by the 3D renderer, a read of a 2x2 texel block for bilinear filtering), having vertically-adjacent elements fall within different memory pages is




to be avoided, as the page crossings required to complete the access typically incur increased memory latencies (and therefore lower performance).

One solution to this problem is to divide the enclosing region into an array of smaller rectangular regions, called memory *tiles*. Surface elements falling within a given tile will all be stored in the same physical memory page, thus eliminating page-crossing penalties for 2D subregion accesses within a tile and thereby increasing performance.

Tiles have a fixed 4KB size and are aligned to physical DRAM page boundaries. They are either 8 rows high by 512 bytes wide or 32 rows high by 128 bytes wide (see *Memory Tile Dimensions*). Note that the dimensions of tiles are irrespective of the data contained within - e.g., a tile can hold twice as many 16-bit pixels (256 pixels/row x 8 rows = 2K pixels) than 32-bit pixels (128 pixels/row x 8 rows = 1K pixels).

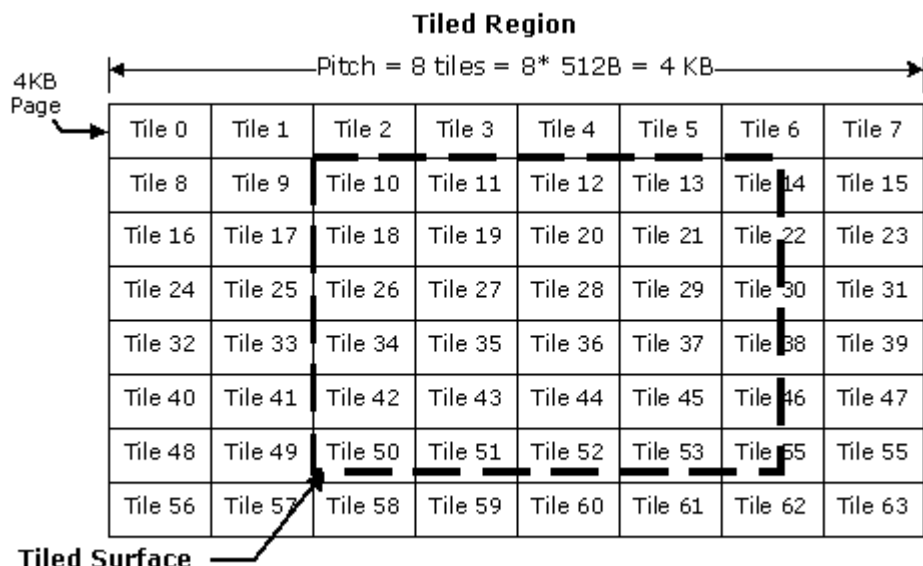
Memory Tile Dimensions



The pitch of a tiled enclosing region must be an integral number of tile widths. The 4KB tiles within a tiled region are stored sequentially in memory in row-major order.

The *Tiled Surface Layout* figure shows an example of a tiled surface located within a tiled region with a pitch of 8 tile widths (512 bytes * 8 = 4KB). Note that it is the *enclosing region* that is divided into tiles - the surface is not necessarily aligned or dimensioned to tile boundaries.

Tiled Surface Layout



B6693-01

Auxiliary Surfaces For Sampled Tiled Resources

For surfaces which are defined as Tiled Resources (TileYs or TileYf format), there may be auxiliary surfaces which are associated with the surface (e.g. HiZ, CCS or MCS). These auxiliary surfaces, while actually not defined as TileYs or TileYf will behave like tiled resources from the hardware perspective. It is possible for software to map and unmap tiles of auxiliary surfaces as tiles of the associated surface are mapped and unmapped. Below is a description how sampling to the mapped/unmapped tile resources is handled for the associated auxiliary surface. Normally, sampling unmapped tiles will return a NULL response to the requesting agen.

For surfaces which are defined as Tiled Resources (TileYs or TileYf format), there may be auxiliary MCS surface which is associated with the surface. These auxiliary surfaces can also be defined as either TileY, TileYs or TileYf. It is possible for software to map and unmap tiles of auxiliary surfaces as tiles of the associated surface are mapped and unmapped. Below is a description how sampling to the mapped/unmapped tile resources is handled for the associated auxiliary surface. Normally, sampling unmapped tiles will return a NULL response to the requesting agen.

MCS

A tile of MCS(Multi-Sample Control Surface) must be mapped to memory whenever MSAA surface pixels associated with the CCS tile are mapped. When all MSAA pixels associated with a MCS tile are unmapped, the MCS may be mapped or unmapped. Below is a table showing the responses for sampling to mapped and unmapped.



Table of Responses for Sampling to MSAA Tiled Resources

MSAA Surface Mapping	MCS Mapping	Sample Response
Mapped	Mapped	Normal Response
Mapped	Unmapped	Undefined Response
Unmapped	Mapped	NULL Response
Unmapped	Unmapped	NULL Response

A "NULL Response" means that the sample returned will be all 0's and the **Null Pixel Mask** (if requested) will indicate the depth pixel is Null.

Tile Formats

Multiple tile formats are supported. The following sections define and describe these formats.

Tiling formats are controlled by programming the fields `Tile_Mode` and `Tiled_Resource_Mode` in the `RENDER_SURFACE_STATE`.

Tile64 Format

This section describes the Tile64 format. Tile64 is chosen with tiling parameters in the Surface State.

Basic shape of 64B in Tile64 is exactly like TileY. In terms of bits per pixels (texels or elements) the following table describes the shape

(width X number of lines or height)

For planar surfaces where the chroma is interleaved, the luma plane and chroma plane start on independent tiles. For NV12 the luma plane bpp is 8bpp while the chroma plane is 16bpp. For P010, the luma plane is 16bpp while chroma is 32bpp.

8bpp	16bpp	32bpp	64bpp	128bpp
16X4	8X4	4X4	2X4	1X4

Contiguous 512B block in the tiled address space consists of 4X2 64B blocks. Therefore the shapes of the pixels (for non-3D/volume and non-MSRTs) are as follow:

8bpp	16bpp	32bpp	64bpp	128bpp
64X8	32X8	16X8	8X8	4X8

Walk within the 4KB tile for 64B's look as the following: (each small rectangle represents 64B)

0	1	2	3	8	9	10	11
4	5	6	7	12	13	14	15
16	17	18	19	24	25	26	27
20	21	22	23	28	29	30	31
32	33	34	35	40	41	42	43
36	37	38	39	44	45	46	47
48	49	50	51	56	57	58	59
52	53	54	55	60	61	62	63

Each 64KB tile shape the following shapes for non-3D/volume and non-MSRTs:

	8bpp	16bpp	32bpp	64bpp	128bpp
64KB in terms of pixels	256X256	256X128	128X128	128X64	64X64
tile array.shape()	2X8	4X4	4X4	8X2	8X2



This layout results in the following (x,y) address bits' swizzle pattern: (from right to left, x, and y bits are from lsbs to msbs)

Bits/Pixel	XY Interleave (TiledAddress[15:0])
8	YYYY YXXY XX YYXXXX
16	YYXX YXXY XX YYXXX-
32	YYXX YXXY XX YYXX--
64	YXXX YXXY XX YYX---
128	YXXX YXXY XX YY----

For 3D/Volume surfaces, each 128B belongs to a unique slice. Therefore, the shapes of pixels X lines X slices for 3D/Volume surfaces are as described below. The stacking represents if the blocks are stacked vertical(V), horizontal(H) or in depth(D) dimensions.

Memory Block	8bpp (X)x(Y)x(Z)	16bpp (X)x(Y)x(Z)	32bpp (X)x(Y)x(Z)	64bpp (X)x(Y)x(Z)	128bpp (X)x(Y)x(Z)
64KB	64X32X32	32X32X32	32X32X16	32X16X16	16X16X16

The actual address bits swizzling in terms of (X,Y,Z) coordinates as follows:

Bits/Pixel	(TiledAddress[15:0])
8	ZZZX YYZY ZX YYXXXX
16	ZZZY YXZY ZX YYXXX-
32	ZZYX YXZY ZX YYXX--
64	ZZYX XXZY ZX YYX---
128	ZZYX XXZY ZX YY----

Further, MSAA planes are also laid out interleaved on two horizontal 64B basis. Therefore, MSAA folding of Q(planeID) looks as shown below. A 64KB tile has up to 4 sample planes. This means that the number of tiles per Q-Slice for X8 is twice and for 16X MSAA is four times as many.

For 8X and 16X IMS multi-sample surfaces, SW must allocate twice the pitch in terms of 64K tiles and for 16X, SW must also allocate 2X the 64K tiles in the y direction.

MSAA Depth/Stencil surface use IMS (Interleaved Multi Samples) which means:

- Use the 1X MSAA (non-MSRT) version of the Tile64 equations and let the client unit do the swizzling internally

- Make sure Depth/Stencil textures have the IMS bit programmed.

so sampler can un-swizzle the samples

Memory Block	8bpp (X)x(Y)x(#pln)	16bpp (X)x(Y)x(#pln)	32bpp (X)x(Y)x(#pln)	64bpp (X)x(Y)x(#pln)	128bpp (X)x(Y)x(#pln)
64KB	256X128X2p	256X64X2p	128X64X2p	128X32X2p	64X32X2p
	128X128X4p	128X64X4p	64X64X4p	64X32X4p	32X32X4p

The actual address swizzling bits in terms of (X,Y,P) coordinates as follows for 2X MSAA and 4X MSAA:

Tile64 : 2X MSAA (P refers to the plane)	
Bits/Pixel	(TiledAddress[15:0])
8	YYYY YYXY PX YXXXXX
16	YYXX YYXY PX YXXXX-
32	YYXX YYXY PX YYXX--
64	YXXX YYXY PX YYX---
128	YXXX YYXY PX YY----

Tile64 : 4X MSAA (P refers to the plane)	
Bits/Pixel	(TiledAddress[15:0])
8	YYYY YYXP PX YXXXXX
16	YYXX YYXP PX YXXXX-
32	YYXX YYXP PX YYXX--
64	YXXX YYXP PX YYX---
128	YXXX YYXP PX YY----

Programming Restriction:

Driver can only re-describe surface to pixel format that is the same size (in terms of total BPP) as the original format.

Re-describing surface to differently sized pixel format will result in image corruption. This is because each BPP has its own tiling formula.

For example: R8G8B8A8_UNORM -> R32_UInt is ok, but R8G8B8A8_UNORM -> R32B32G32A32_UInt will result in corruption.

Following section describes the Tile64 Layout Change for 3D/Volume surface and planar multi-sample surfaces.



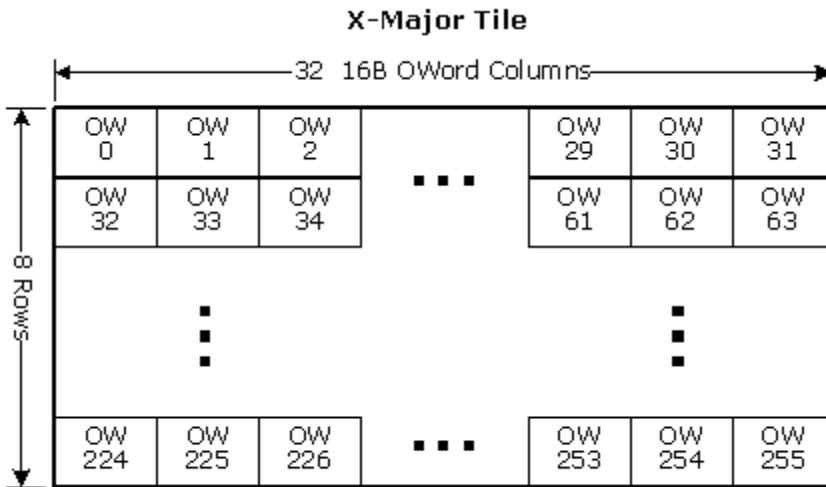
Tile-X Legacy Format

The legacy format Tile-X is a *X-Major* (row-major) storage of tile data units, as shown in the following figure. It is a 4KB tile which is subdivided into an 8-high by 32-wide array of 16-byte OWords . The selection of tile direction only impacts the internal organization of tile data, and does not affect how surfaces map onto tiles. Note that an X-major tiled region with a tile pitch of 1 tile is actually stored in a linear fashion.

Tile-X format is selected for a surface by programming the Tiled_Mode field in RENDER_SURFACE_STATE to XMAJOR.

For 3D sampling operation, a surface using Tile-X layout is generally lower performance the organization of texels in memory.

Tile X-Tile (X-Major) Layout



B6694-01

Tile4 Format

This section describes the definition of Tile4.

Tile4 is chosen with tiling parameters in the Render Surface State.

Basic shape of 64B in Tile4 is exactly like the previously supported TileY.

A walk within the 4KB tile for 64B's look as the following: (each small rectangle represents 64B)

0	1	2	3	8	9	10	11
4	5	6	7	12	13	14	15
16	17	18	19	24	25	26	27
20	21	22	23	28	29	30	31
32	33	34	35	40	41	42	43
36	37	38	39	44	45	46	47
48	49	50	51	56	57	58	59
52	53	54	55	60	61	62	63

Layout of 3D/Volume surfaces using Tile4 format continue to be similar to the legacy TileY format. Therefore, beyond 4K tile, there is no change for this format for these special layouts.

Tile4 does not support MIP Tail. Tile4 also does not support MSAA.



Tiling Algorithm

This section describes the pseudo-code for redefined TileY and TileYs. In the pseudo-code, Tile4 refers to 4K shape of a tile which is TileY format and Tile64 refers to 64K shape of the tile.

inputs (clx,cly) is cartesian coordinates for a CL, P = pitch in terms of tiles, BPP = bytes per pixel

BA = base address of allocation R2T_Xoffset, R2T_Yoffset

tile4W = 2 # number of CL_blocks(512B) across in 4K tile

tile4H = 4 # number of lines down in the 4K tile

cBlkW = 4 # number of CLs in X dim within CLblock of 512B

cBlkH = 2 # number of CLs in Y dim within CLblock of 512B

x = clx + R2T_x_offset

y = cly + R2T_y_offset

bpp = {0,1,2,3,4} # 8bpp, 16bpp, 32bpp, 64bpp, 128bpp

tile64 specific

tile64W = bpp? 8:8:4:4:2 # 2 for 8bpp

tile64H = bpp? 2:2:4:4:8 # 8 for 8bpp

tile64_size = tile64W*tile64H*tile4_size

tile64_rowsize = p*tile64_size # tile row size in terms of 64B cachelines

num_tile64_in_x = x div (tile64W*tile4W*cBlkW)

clx_in_tile64 = x % (tile64W*tile4W*cBlkW)

num_tile64_in_y = y div (tile64H*tile4H*cBlkH)

cly_in_tile64 = y % (tile64H*tile4H*cBlkH)

if (tile_format == tile4):

x_muxed = x

y_muxed = y

else:

x_muxed = clx_in_tile64

y_muxed = cly_in_tile64

#####

tile4_size = tile4W*tile4H*cBlockW*cBlockH

tile4_rowsize = p*tile4_size # tile row size in terms of 64B cachelines

num_tile4_in_x = x_muxed div (tile4W*cBlkW)

```

clx_in_tile4 = x_muxed % (tile4W*cBlkW)
num_tile4_in_y = y_muxed div (tile4H*cBlkH)
cly_in_tile4 = y_muxed % (tile4H*cBlkH)
num_blk_in_x = clx_in_tile4 div cBlkW
clx_in_blk = clx_in_tile4 % cBlkW
num_blk_in_y = cly_in_tile4 div cBlkH
cly_in_blk = cly_in_tile4 % cBlkH
tiledOffset_tile4 = num_tile4_in_y * tile4_rowsize +
num_tile4_in_x * tile4_size +
num_blk_in_y * tile4W * cBlkW * cBlkH +
num_blk_in_x * cBlkW * cBlkH +
cly_in_blk * cBlkH +
clx_in_blk
tiledOffset_tile64 = num_tile64_in_y * tile64_rowsize +
num_tile64S_in_x * tile64_size +
num_tile64_in_y * tile64W * tile64_size +
num_tile64_in_x * tile64_size +
num_blk_in_y * tile4W * cBlkW * cBlkH +
num_blk_in_x * cBlkW * cBlkH +
cly_in_blk * cBlkH +
clx_in_blk
if (tile_format == tile4):
tiledOffset = tiledOffset_tile4
else:
tiledOffset = tiledOffset_tile64
TileVirtualAddress = BA + tiledOffset

```



Tiling Support

The rearrangement of the surface elements in memory must be accounted for in device functions operating upon tiled surfaces. (Note that not all device functions that access memory support tiled formats). This requires either the modification of an element's linear memory address or an alternate formula to convert an element's X,Y coordinates into a tiled memory address.

However, before tiled-address-generation can take place, some mechanism must be used to determine whether the surface elements accessed fall in a linear or tiled region of memory, and if tiled, what the tile region pitch is, and whether the tiled region uses X-Major or Y-Major format. There are two mechanisms by which this detection takes place: (a) an implicit method by detecting that the pre-tiled (linear) address falls within a "fenced" tiled region, or (b) by an explicit specification of tiling parameters for surface operands (i.e., parameters included in surface-defining instructions).

The following table identifies the tiling-detection mechanisms that are supported by the various memory streams.

Access Path	Tiling-Detection Mechanisms Supported
Processor access through the Graphics Memory Aperture	Fenced Regions
3D Render (Color/Depth Buffer access)	Explicit Surface Parameters
Sampled Surfaces	Explicit Surface Parameters
Blt operands	Explicit Surface Parameters
Display and Overlay Surfaces	Explicit Surface Parameters

Tiled (Fenced) Regions

The only mechanism to support the access of surfaces in tiled format by the host or external graphics client is to place them within "fenced" tiled regions within Graphics Memory. A fenced region is a block of Graphics Memory specified using one of the sixteen FENCE device registers. (See *Memory Interface Registers* for details). Surfaces contained within a fenced region are considered tiled from an external access point of view. Note that fences cannot be used to untile surfaces in the PGM_Address space since external devices cannot access PGM_Address space. Even if these surfaces (or any surfaces accessed by an internal graphics client) fall within a region covered by an enabled fence register, that enable will be effectively masked during the internal graphics client access. Only the explicit surface parameters described in the next section can be used to tile surfaces being accessed by the internal graphics clients.

Tiled Surface Parameters

Internal device functions require explicit specification of surface tiling parameters via information passed in commands and state. This capability is provided to limit the reliance on the fixed number of fence regions.

The following table lists the surface tiling parameters that can be specified for 3D Render surfaces (Color Buffer, Depth Buffer, Textures, etc.) via SURFACE_STATE.

Surface Parameter	Description
Tiled Surface	If ENABLED, the surface is stored in a tiled format. If DISABLED, the surface is stored in a linear format.
Tile Walk	If Tiled Surface is ENABLED, this parameter specifies whether the tiled surface is stored in Y-Major or X-Major tile format.
Base Address	Additional restrictions apply to the base address of a Tiled Surface vs. that of a linear surface.
Pitch	Pitch of the surface. Note that, if the surface is tiled, this pitch must be a multiple of the tile width.

Tiled Surface Restrictions


Additional restrictions apply to the Base Address and Pitch of a surface that is tiled. In addition, restrictions for tiling via SURFACE_STATE are subtly different from those for tiling via fence regions. The most restricted surfaces are those that will be accessed both by the host (via fence) and by internal device functions. An example of such a surface is a tiled texture that is initialized by the CPU and then sampled by the device.

The tiling algorithm for internal device functions is different from that of fence regions. Internal device functions always specify tiling in terms of a surface. The surface must have a base address, and this base address is not subject to the tiling algorithm. Only *offsets* from the base address (as calculated by X, Y addressing within the surface) are transformed through tiling. The base address of the surface must therefore be 4KB-aligned. This forces the 4KB tiles of the tiling algorithm to exactly align with 4KB device pages once the tiling algorithm has been applied to the offset. The width of a surface must be less than or equal to the surface pitch. There are additional considerations for surfaces that are also accessed by the host (via a fence region).

Fence regions have no base address per se. Host linear addresses that fall in a fence region are translated in their entirety by the tiling algorithm. It is as if the surface being tiled by the fence region has a base address in graphics memory equal to the fence base address, and all accesses of the surfaces are (possibly quite large) offsets from the fence base address. Fence regions have a virtual "left edge" aligned with the fence base address, and a "right edge" that results from adding the fence pitch to the "left edge". Surfaces in the fence region must not straddle these boundaries.

Base addresses of surfaces that are to be accessed both by an internal graphics client and by the host have the tightest restrictions. In order for the surface to be accessed without GTT re-mapping, the surface base address (as set in SURFACE_STATE) must be a "Tile Row Start Address" (TRSA). The first address in each tile row of the fence region is a Tile Row Start Address. The first TRSA is the fence base address. Each TRSA can be generated by adding an integral multiple of the row size to the fence base address. The row size is simply the fence pitch in tiles multiplied by 4KB (the size of a tile.)

Tiled Surface Placement



The pitch in SURFACE_STATE must be set equal to the pitch of the fence that will be used by the host to access the surface if the same GTT mapping will be used for each access. If the pitches differ, a different GTT mapping must be used to eliminate the "extra" tiles (4KB memory pages) that exist in the excess rows at the right side of the larger pitch. Obviously no part of the surface that will be accessed can lie in pages that exist only in one mapping but not the other. The new GTT mapping can be done manually by SW between the time the host writes the surface and the device reads it, or it can be accomplished by arranging for the client to use a different GTT than the host (the PPGTT -- see *Logical Memory Mapping* below).

The width of the surface (as set in SURFACE_STATE) must be less than or equal to both the surface pitch and the fence pitch in any scenario where a surface will be accessed by both the host and an internal graphics client. Changing the GTT mapping will not help if this restriction is violated.

Surface Access	Base Address	Pitch	Width	Tile "Walk"
Host only	No restriction	Integral multiple of tile size <= 256KB	Must be <= Fence Pitch	No restriction
Client only	4KB-aligned	Integral multiple of tile size <= 256KB	Must be <= Surface Pitch	Restrictions imposed by the client (see Per Stream Tile Format Support)
Host and Client, No GTT Remapping	Must be TRSA	Fence Pitch = Surface Pitch = integral multiple of tile size <= 256KB	Width <= Pitch	Surface Walk must meet client restriction, Fence Walk = Surface Walk
Host and Client, GTT Remapping	4KB-aligned for client (will be tile-aligned for host)	Both must be Integral multiple of tile size <=128KB, but not necessarily the same	Width <= Min(Surface Pitch, Fence Pitch)	Surface Walk must meet client restriction, Fence Walk = Surface Walk



Per-Stream Tile Format Support

MI Client	Tile Formats Supported	
CPU Read/Write	All	
Display/Overlay	Y-Major not supported X-Major required for Async Flips	
Blt	Linear and X-Major only No Y-Major support	
3D Sampler	All Combinations of TileY, TileX and Linear are supported. TileY is the fastest, Linear is the slowest.	
3D Color,Depth	Rendering Mode Color-vs-Depth bpp	Buffer Tiling Supported
	Classical Same Bpp	Both Linear Both Tile4 and Tile64 Both TileY Linear & TileX Linear & (Tile4 or Tile64) TileX & (Tile4 or Tile64)
	Classical Mixed Bpp	Both Linear Both TileX Both Tile4 and Tile64 Linear & TileX Linear & (Tile4 or Tile64) TileX & (Tile4 or Tile64)

Memory Compression

CCS Surface Encodings

This page presents CCS encodings for projects prior to Unified Compression feature.

The compression status is held as a 4 bit encoding in Compression Control Surface. For 3D one 4bit encoding represents status of 128B of data.

For Media, we use 3rd party compressor capable of handling 256b block compression, 4bit encoding is duplicated for left and right 128B.

Table of CCS encodings

Color Code	Notes
x	Unused data
clr	Data at clear, needs to be filled in with clear value, no actual data is transferred
un	Uncompressed data
cp	Compressed data

Unified Compression, Compression Control Surface encodings																	
		MEDIA								3D							
CCS		Media	Notes	CL3 (64B)		CL2(64B)		CL1(64B)		CL0(64B)		3D	Notes	CL1(64B)		CL0(64B)	
0000	0x0	Media_1to1	Uncompressed	u	u	u	u	u	u	u	u	3D_1to1	Uncompressed	u	u	u	u
0001	0x1	Media_4to1_Mono	128B to 32B					x	x	x	cp	3D_4to1_Mono	128B to 32B	x	x	x	cp
0010	0x2	Media_4to3_UCU	not used, 128B, upper uncompressed					u	u	x	cp	3D_4to3_UCU	upper uncompressed, lower comp	u	u	x	cp
0011	0x3	Media_8to6_UCU	not used, 256B, upper uncompressed	u	u	u	u	x	x	cp	cp	3D_2to1_CL	upper uncompressed and lower	u	u	clr	clr

													clear					
0100	0x4	Media_8to1_Mono	256B to 32B shared	x	x	x	x	x	x	x	c	p	3D_FDVO	FDVO entire cache line filled with zero	x	x	x	x
0101	0x5	Media_8to2_Mono	256B to 64B shared	x	x	x	x	x	x	c	p	c	3D_Null_Tile	Tile is Null	x	x	x	x
0110	0x6	Media_4to2_Mono	128B to 64B					x	x	c	p	c	3D_4to2_Mono	128B to 64B	x	x	c	p
0111	0x7	Media_8to6_UCL	not used, 256B, lower uncompressed	x	x	c	p	c	p	u	n	u	N/A	not used	x	x		x
1000	0x8	Media_4to3_UCL	not used, 128B, lower uncompressed					x	c	p	u	n	3D_4to3_UCL	lower uncompressed	x	c	u	n
1001	0x9	Media_4to3_Mono	not used, monolithic 4:3 Feature					x	c	p	c	p	3D_4to3_Mono	not used, monolithic 4:3 Feature	x	c	c	p
1010	0xA	Media_4to2_Individual	not used, 128B, 64B individually compressed					x	c	p	x	c	3D_4to2_Individual	64B individually compressed	x	c	x	c
1011	0xB	Media_8to3_Mono	256B to 96B	x	x	x	x	x	c	p	c	p	3D_4to1_CCL	lower clear, upper compressed	x	c	c	cl
1100	0xC	Media_8to4_Mono	256B to 128B	x	x	x	x	c	p	c	p	c	3D_2to1_CCU	upper clear and lower uncompressed	cl	cl	u	n

													essed				
11 10	0x E	Media_8to6_ Mono	256B to 192B	x	x	c p	c p	c p	c p	c p	c p	3D_4to1_C CU	upper clear, lower compres sed	cl r	cl r	x	c p
11 11	0x F	Media_8to7_ Mono	not used, 256B to224B shared	x	c p	c p	c p	c p	c p	c p	c p	3D_ClearCl ear	Both 64B chunks at clear value	cl r	cl r	c lr	cl r

Media Memory Compression

The software requirement when using media memory compression is to allocate each compressible surface one memory tile wider than is required based on the surface width plus normal byte padding (this approach is called "pitch+1"). The reason is each compressible surface needs an "extra" tile to the right edge of surface to store important compression control information. For example, if the surface is 1920x1088, this would normally be allocated by the driver to be 2048 bytes wide, or 16 tiles (for NV12 8bpp). Using this "pitch + 1", the pitch would be set to 17 instead of 16 (and the surface width remains unchanged, only pitch is increased).

The largest supported width will be 4K pixels for 2D RGBA 8bpp surfaces and 2x2K for S3D surfaces (for 4KB pages). E.g. the pitch would be set to 129 in these cases (128+1). NV12 4K would be 33 (28+1). The case of 64KB pages is the same: the driver will allocate 1 extra page to the right ("pitch + 1"), however now the 4K wide restriction is relaxed. With 64KB pages, the widest surface that supports memory compression is 16K for 2D RGBA 8bpp or 2x8K for S3D. E.g. the pitch would be set to 129 in these cases (128+1).

Memory Object Overview

Any memory data accessed by the device is considered part of a *memory object* of some memory object type.

The following table lists the various memory objects types and an indication of their role in the system.

Memory Object Type	Role
Graphics Translation Table (GTT)	Contains PTEs used to translate "graphics addresses" into physical memory addresses.
Hardware Status Page	Cached page of systemem used to provide fast driver synchronization.
Logical Context Buffer	Memory areas used to store (save/restore) images of hardware rendering contexts. Logical contexts are referenced via a pointer to the corresponding Logical Context Buffer.
Ring Buffers	Buffers used to transfer (DMA) instruction data to the device. Primary means of controlling rendering operations.

Memory Object Type	Role
Batch Buffers	Buffers of instructions invoked indirectly from Ring Buffers.
State Descriptors	Contains state information in a prescribed layout format to be read by hardware. Many different state descriptor formats are supported.
Vertex Buffers	Buffers of 3D vertex data indirectly referenced through "indexed" 3D primitive instructions.
VGA Buffer (Must be mapped UC on PCI)	Graphics memory buffer used to drive the display output while in legacy VGA mode.
Display Surface	Memory buffer used to display images on display devices.
Overlay Surface	Memory buffer used to display overlaid images on display devices.
Overlay Register, Filter Coefficients	Memory area used to provide double-buffer for Overlay register and filter coefficient loading.
Cursor Surface	Hardware cursor pattern in memory.
2D Render Source	Surface used as primary input to 2D rendering operations.
2D Render R-M-W Destination	2D rendering output surface that is read in order to be combined in the rendering function. Destination surfaces that accessed via this Read-Modify-Write mode have somewhat different restrictions than Write-Only Destination surfaces.
2D Render Write-Only Destination	2D rendering output surface that is written but not read by the 2D rendering function. Destination surfaces that accessed via a Write-Only mode have somewhat different restrictions than Read-Modify-Write Destination surfaces.
2D Monochrome Source	1 bpp surfaces used as inputs to 2D rendering after being converted to foreground/background colors.
2D Color Pattern	8x8 pixel array used to supply the "pattern" input to 2D rendering functions.
DIB	"Device Independent Bitmap" surface containing "logical" pixel values that are converted (via LUTs) to physical colors.
3D Color Buffer	Surface receiving color output of 3D rendering operations. May also be accessed via R-M-W (aka blending). Also referred to as a Render Target.
3D Depth Buffer	Surface used to hold per-pixel depth and stencil values used in 3D rendering operations. Accessed via RMW.
3D Texture Map	Color surface (or collection of surfaces) which provide texture data in 3D rendering operations.
"Non-3D" Texture	Surface read by Texture Samplers, though not in normal 3D rendering operations (for example, in video color conversion functions).
Motion Comp Surfaces	These are the Motion Comp reference pictures.
Motion Comp Correction Data Buffer	This is Motion Comp intra-coded or inter-coded correction data.