Intel® Iris® Xe and UHD Graphics Open Source

Programmer's Reference Manual

For the 2020-2021 11th Generation Intel Xeon®, Core™, Celeron®, Pentium® Gold Processors based on the "Tiger Lake" Platform

Volume 4: Configurations

December 2021, 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.

Table of Contents

Configurations

This chapter contains configurations details as described in the following sections:

- Top Level Block Diagrams
- Device Attributes
- Steppings and Device IDs

Product Mapping Table

Product Configuration Attribute Table			
Product Family	**TGL**		
SKU Name	**UP4**	**UP3**	**H81**
Status	**POR**	**POR**	**POR**
Global Attributes			
Render Engine	1x6x16	1x6x16	1x2x16
Media Engine	Xe_M	Xe_M	Xe_M
Display Engine	Xe_D	Xe_D	Xe_D
Chassis			
LLC Size	12MB	12MB	24MB
In-Package Memory	N/A	N/A	N/A

Top Level Block Diagrams

Overview

The slice is constructed of 6 dual subslices, each containing 16eus and a double sampler each capable of 8 tex/clk. Slice common will be constructed of a single Z pipe, capable of 16x8 HIZ and 8x4 IZ. 3 Color pipes, each of 8f8b; sharing the compression block for area efficiency. Total L3 cache size is defined to be 4MB, built from 8 banks of 480kB each. Initial sizing includes 50% growth of shared L3 data cache and URB, from 256kB->384kB per bank and maintains 128kB of tile cache per bank.

The machine will consist of a single geometry fixed function pipeline and maintain the position only shading (POSH) pipeline introduced in previous IP.

It will include support for a second IDI port to sustain the ~100GB/s memory bandwidth requirement to support the 96eu machine. Memory subsystem in this timeframe expected to be of similar bandwidth. As such, a single IDI is no longer sufficient as peak BW is limited to 64GB/s per direction at 1GHz. A single GAM will be designed to support both IDIs, offering a combined 128r+128w bytes per clock.

Media fixed function blocks are as follows: 2 VDBoxes, 1 VEBox, and 2 SFC. These assets are deemed sufficient to meet usage and throughput requirements for a majority of client SKUs. Within the slice, there will be a total of 6 VME and 6 AVS; built 1 per dual subslice. While there are inversions on VME assets vs previous generations; native support for HEVC 422 encode and random access will enable VDEnc to serve these usages at a lower power than previously offered with VME.

1x6x16

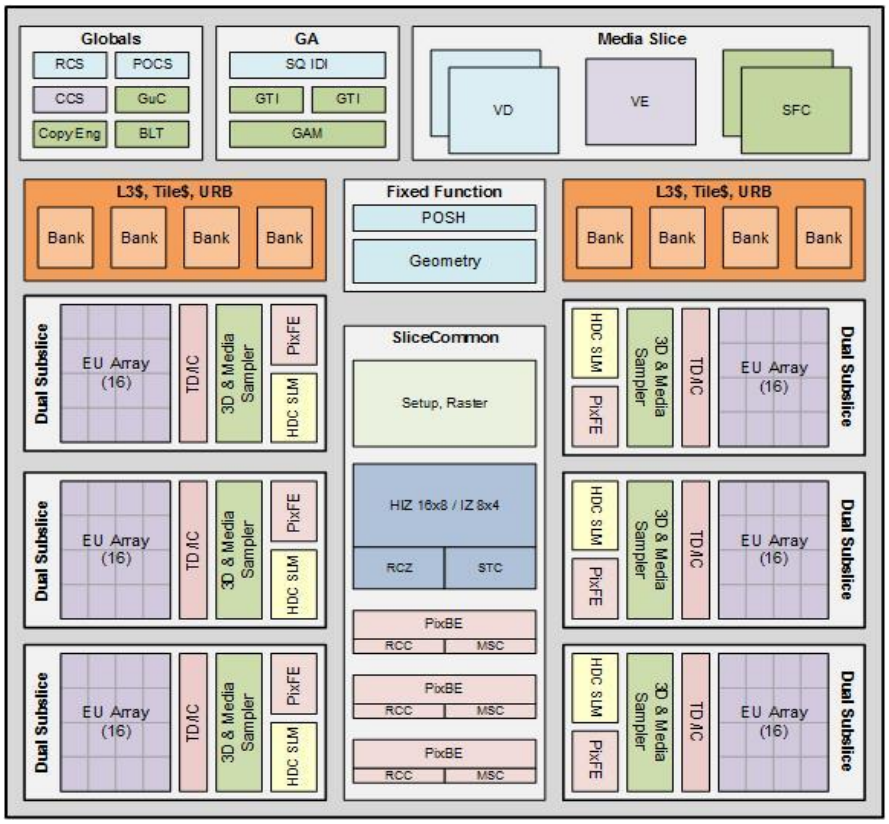

Device Attributes

The following table lists detailed GT device attributes for proposed SKUs.

NOTE: This information is preliminary, and subject to change.

Product Family	TGL			
Architectural Name	1x2x16	1x3x16	1x5x16	1x6x16
SKU Name	GT1	GT1.xf	GT1.xf	GT2
	Global Attributes			
Slice Count	1	1	1	1
Dual-Subslice (DSS) Count	2	3	5	6
EU/DSS	16	16	16	16
EU Count (Total)	32	48	80	96
Threads / EU	7	7	7	7
Thread Count (Total)	224	336	560	672
FLOPs/Clk - Half Precision, MAD (peak)	1024	1536	2560	3072
FLOPs/Clk - Single Precision, MAD (peak)	512	768	1280	1536
FLOPs/Clk - Double Precision, MAD (peak)	N/A	N/A	N/A	N/A
Unslice clocking (coupled/decoupled from Cr slice)	Coupled	Coupled	Coupled	Coupled
GTI/Ring Interfaces	2	2	2	2
GTI bandwidth (bytes/unslice-clk)	r: 128 w: 128	r: 128 w: 128	r: 128 w: 128	r: 128 w: 128
eDRAM Support	N/A	N/A	N/A	N/A
Graphics Virtual Address Range	48 bit	48 bit	48 bit	48 bit
Graphics Physical Address Range	39 bit	39 bit	39 bit	39 bit
	Cacches & Dedicated Memories			
L3 Cache, total size (bytes)[1]	1920k	2880k	3840k	3840k
L3 Cache, bank count[1]	4	6	8	8
L3 Cache, bandwidth (bytes/clk)	4x 64 R|W	6x 64 R|W	8x 64 R|W	8x 64 R|W
L3 Cache, URB bandwidth (bytes/clk)[(1)]	4x 64 R|W	6x 64 R|W	8x 64 R|W	8x 64 R|W
L3 Cache, D$ Size (Kbytes)[(2)]	704K	1056K	1408K	1408K
L3 Cache, Tile cache size (Kbytes)[(2)]	1024K	1536K	2048K	2048K
L3 Cache, Command buffer cache size (Kbytes)[(2)]	64K	96K	128K	128K
URB Size (kbytes)[(2)]	256K	384K	512K	512K
SLM Size (kbytes)	256k	384k	640k	768k
Instruction Cache (instances, bytes ea.)	2x 48k	3x 48k	5x 48k	6x 48k

Color Cache (RCC, bytes)	1x 32k	2x 32k	3x 32k	3x 32k
MSC Cache (MSC, bytes)	1x 16k	2x 16k	3x 16k	3x 16k
HiZ Cache (HZC, bytes)	1x 12k	1x 12k	1x 12k	1x 12k
Z Cache (RCZ, bytes)	1x 32k	1x 32k	1x 32k	1x 32k
Stencil Cache (STC, bytes)	1x 8k	1x 8k	1x 8k	1x 8k
	Instruction Issue Rates			
FMAD, SP (ops/EU/clk)	8	8	8	8
FMUL, SP (ops/EU/clk)	8	8	8	8
FADD, SP (ops/EU/clk)	8	8	8	8
MIN,MAX, SP (ops/EU/clk)	8	8	8	8
CMP, SP (ops/EU/clk)	8	8	8	8
INV, SP (ops/EU/clk)	2	2	2	2
SQRT, SP (ops/EU/clk)	2	2	2	2
RSQRT, SP (ops/EU/clk)	2	2	2	2
LOG, SP (ops/EU/clk)	2	2	2	2
EXP, SP (ops/EU/clk)	2	2	2	2
IDIV, SP (ops/EU/clk)	1-6	1-6	1-6	1-6
TRIG, SP (ops/EU/clk)	2	2	2	2
	Load Store			
Data Ports (HDC)	2	3	5	6
L3 Load/Store - same addresses within msg (Bytes/clk)	128	192	320	384
L3 Load/Store - unique addresses within msg (Bytes/clk)	128	192	320	384
SLM Load//Store - same addresses within msg (Bytes/clk)	256	384	640	768
SLM Load//Store - unique addresses within msg (Bytes/clk)	256	384	640	768
Atomic, Local 32b - same addresses within msg (dwords/clk)	2	3	5	6
Atomic, Global 32b - unique addresses within msg (dwords/clk)	32	48	80	96
	3D Attributes			
Geometry pipes	1	1	1	1
Samplers (3D)	2	3	5	6
2D Texel Rate, point, 32b (tex/clk)	16	24	40	48
2D Texel Rate, point, 64b (tex/clk)	16	24	40	48
2D Texel Rate, point, 128b (tex/clk)	16	24	40	48
2D Texel Rate, bilinear, 32b (tex/clk)	16	24	40	48
2D Texel Rate, bilinear, 64b (tex/clk)	16	24	40	48
2D Texel Rate, bilinear, 128b (tex/clk)	4	6	10	12

2D Texel Rate, trilinear, 32b (tex/clk)	8	12	20	24
2D Texel Rate, trilinear, 64b (tex/clk)	8	12	20	24
2D Texel Rate, trilinear, 128b (tex/clk)	2	3	5	6
2D Texel Sample Rate, aniso 2x (MIP nearest), 32b (tex/clk)	16	24	40	48
2D Texel Sample Rate, aniso 4x (MIP nearest), 32b (tex/clk)	8	12	20	24
2D Texel Sample Rate, aniso 8x (MIP nearest), 32b (tex/clk)	4	6	10	12
2D Texel Sample Rate, aniso 16x (MIP nearest), 32b (tex/clk)	2	3	5	6
3D Texel Sample Rate, point, 32b (tex/clk)	16	24	40	48
3D Texel Sample Rate, point, 64b (tex/clk)	16	24	40	48
3D Texel Sample Rate, point, 128b (tex/clk)	8	12	20	24
3D Texel Sample Rate, bilinear, 32b (tex/clk)	8	12	20	24
3D Texel Sample Rate, bilinear, 64b (tex/clk)	8	12	20	24
3D Texel Sample Rate, bilinear, 128b (tex/clk)	2	3	5	6
HiZ Rate, (ppc)	1x128	1x128	1x128	1x128
IZ Rate, (ppc)	1x32	1x32	1x32	1x32
Stencil Rate (ppc)	1x128	1x128	1x128	1x128
(500 MHz, DDR-4267; Range depends on dynamic compression ratio)				
Pixel Rate, fill, 32bpp (pix/clk, RCC hit)	8	12	20	24
Pixel Rate, fill, 32bpp (pix/clk, LLC hit @ 1.0x unslice clk)	8	12	20	24
Pixel Rate, fill, 32bpp (pix/clk, LLC hit, @ 1.5x unslice clk)	N/A	N/A	N/A	N/A
Pixel Rate, fill, 32bpp (pix/clk, memory, @ 1.0x unslice clk)	8	12	20	24
Pixel Rate, fill, 32bpp (pix/clk, memory, @ 1.5x unslice clk)	N/A	N/A	N/A	N/A
(500 MHz, DDR-4267 or eDRAM; Range depends on dynamic compression ratio)				
Pixel Rate, blend, 32bpp (p/clk, RCC hit)	8	12	20	24
Pixel Rate, blend, 32bpp (p/clk, RCC miss, @ 1.0x unslice clk)	8	12	20	24
Pixel Rate, blend, 32bpp (p/clk, RCC miss, @ 1.5x unslice clk)	N/A	N/A	N/A	N/A
Pixel Rate, blend, 32bpp (pix/clk, memory, @ 1.0x unslice clk)	8	12	17-20	17-24
Pixel Rate, blend, 32bpp (pix/clk, memory, @ 1.5x unslice clk)	N/A	N/A	N/A	N/A

Notes:

* Architectural Name = Slice Count x Subslice Count x EUs per Subslice

(1) L3 cache and URB share the write bandwidth. Read bandwidths of 64B/clk can be achieved independently

(2) URB/Data cache/Tile cache/Command buffer cache sizes are programmable. Indicative values presented in this table.

Device Attributes Xe_M

NOTE: This information is preliminary, and subject to change.

	Product Configuration Attribute Table	
Product Family	**TGL**	
Architectural Name	**1x2x16**	**1x6x16**
SKU Name	**GT1**	**GT2**
	Media Attributes	
Samplers (VME)	2	6
Samplers (AVS)	2	6
VDBox Instances (See)	2	2
VEBox Instances	1	1
SFC Instances	2	2

Device Attributes Display

Refer to display overview

Steppings and Device IDs

SKUs and Device IDs

The following table details all currently planned SKUs. This information is subject to change at any time based on roadmap plans or new steppings.

Segment	SKU	Config	EUs	VDBoxes	TDP	Gfx Name	CPU Stepping	GT Version	Display Version	Dev2 ID	RevID
Mobile	UP3 / H35	1x6x16	96,80	2	15W / 35W	Intel(R) Iris(R) Xe Graphics	B2	GT2 B0	C0	0x9A49	0x1
Mobile	UP3	1x3x16	48	1	15W	Intel(R) UHD Graphics	B2	GT2 B0	C0	0x9A78	0x1
Mobile	UP3 / H35	1x6x16	96,80	2	15W / 35W	Intel(R) Iris(R) Xe Graphics	C0	GT2 C0	D0	0x9A49	0x3
Mobile	UP3	1x3x16	48	1	15W	Intel(R) UHD Graphics	C0	GT2 C0	D0	0x9A78	0x3
Mobile	UP4	1x6x16	96,80	2	5.2W	Intel(R) Iris(R) Xe Graphics	B2	GT2 B0	C0	0x9A40	0x1
Mobile	UP4	1x3x16	48	1	5.2W	Intel(R) UHD Graphics	B2	GT2 B0	C0	0x9A78	0x1
Mobile	UP4	1x6x16	96,80	2	5.2W	Intel(R) Iris(R) Xe Graphics	C0	GT2 C0	D0	0x9A40	0x3
Mobile	UP4	1x3x16	48	1	5.2W	Intel(R) UHD Graphics	C0	GT2 C0	D0	0x9A78	0x3
Perf Mobile	H81	1x2x16	32	2	45W	Intel(R) UHD Graphics	R0	GT1 B0	D0	0x9A60	0x1
Perf Mobile Low Graphics	H81	1x1x16	16	1	45W	Intel(R) UHD Graphics	R0	GT1 B0	D0	0x9A68	0x1
Mobile WS	H81	1x2x16	32	2	45W	Intel(R) UHD Graphics	R0	GT1 B0	D0	0x9A70	0x1

Notes:

The Intel UHD Graphics Device ID SKUs are unified for both UP3 and UP4, eg: there is no unique device ID between UP3 and UP4 for 48 EU devices.

The UP3 and H35 SKUs share the same DevIDs and are equivalent other than being focused on different TDPs.