

10th Generation Intel[®] Core[™] Processor Families

Datasheet, Volume 1 of 2

Supporting 10th Generation Intel[®] Core[™] Processor Families, Intel[®] Pentium[®] Processors, Intel[®] Celeron[®] Processors for U/Y Platforms, formerly known as Ice Lake

July 2020

Revision 005



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.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. **No computer system can be absolutely secure.** Check with your system manufacturer or retailer or learn more at intel.com.

Intel technologies may require enabled hardware, specific software, or services activation. Check with your system manufacturer or retailer.

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.

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.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visit www.intel.com/design/literature.htm. No computer system can be absolutely secure.

Intel, Core, Pentium, VTune, Thunderbolt, and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Copyright © 2019-2020, Intel Corporation. All rights reserved.




Contents

1	Introduction	12
1.1	Processor Volatility Statement	13
1.2	Package Support	13
1.3	Supported Technologies	13
1.3.1	API Support (Windows*)	14
1.4	Power Management Support	14
1.4.1	Processor Core Power Management	14
1.4.2	System Power Management	14
1.4.3	Memory Controller Power Management	14
1.4.4	Processor Graphics Power Management	15
1.4.4.1	Memory Power Savings Technologies	15
1.4.4.2	Display Power Savings Technologies	15
1.4.4.3	Graphics Core Power Savings Technologies	15
1.5	Thermal Management Support	15
1.6	Processor Testability	16
1.7	Operating Systems Support	16
1.8	Terminology and Special Marks	16
2	Technologies	20
2.1	Platform Environmental Control Interface (PECI)	20
2.1.1	PECI Bus Architecture	20
2.2	Intel® Virtualization Technology (Intel® VT)	22
2.2.1	Intel® Virtualization Technology (Intel® VT) for Intel® 64 and Intel® Architecture (Intel® VT-X)	22
2.2.2	Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d)	23
2.2.3	Intel® APIC Virtualization Technology (Intel® APICv)	26
2.3	Security Technologies	27
2.3.1	Intel® Trusted Execution Technology (Intel® TXT)	27
2.3.2	Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)	28
2.3.3	Perform Carry-Less Multiplication Quad Word (PCLMULQDQ) Instruction	28
2.3.4	Intel® Secure Key	28
2.3.5	Execute Disable Bit	29
2.3.6	Boot Guard Technology	29
2.3.7	Intel® Supervisor Mode Execution Protection (SMEP)	29
2.3.8	Intel® Supervisor Mode Access Protection (SMAP)	30
2.3.9	Intel® Software Guard Extensions (Intel® SGX)	30
2.3.10	Intel® Secure Hash Algorithm Extensions (Intel® SHA Extensions)	31
2.3.11	User Mode Instruction Prevention (UMIP)	31
2.3.12	Read Processor ID (RDPID)	32
2.4	Power and Performance Technologies	32
2.4.1	Intel® Smart Cache Technology	32
2.4.2	IA Core Level 1 and Level 2 Caches	32
2.4.3	Intel® Turbo Boost Max Technology 3.0	33
2.4.4	Power Aware Interrupt Routing (PAIR)	34
2.4.5	Intel® Hyper-Threading Technology (Intel® HT Technology)	34
2.4.6	Intel® Turbo Boost Technology 2.0	34
2.4.6.1	Intel® Turbo Boost Technology 2.0 Power Monitoring	34
2.4.6.2	Intel® Turbo Boost Technology 2.0 Power Control	35
2.4.6.3	Intel® Turbo Boost Technology 2.0 Frequency	35
2.4.7	Enhanced Intel SpeedStep® Technology	35
2.4.8	Intel® Speed Shift Technology	36




2.4.9	Intel® Advanced Vector Extensions 2 (Intel® AVX2)	36
2.4.10	Intel® 64 Architecture x2APIC	36
2.4.11	Intel® GNA (GMM and Neural Network Accelerator)	37
2.4.12	Advanced Vector Extensions 512 Bit (Intel® AVX-512)	38
2.4.13	Cache Line Write Back (CLWB)	39
2.4.14	Ring Interconnect	39
2.5	Intel® Image Processing Unit (Intel® IPU)	40
2.5.1	Platform Imaging Infrastructure	40
2.5.2	Intel® Image Processing Unit (Intel® IPU)	40
2.6	Debug Technologies	41
2.6.1	Intel® Processor Trace	41
3	Power Management	42
3.1	Advanced Configuration and Power Interface (ACPI) States Supported	43
3.2	Processor IA Core Power Management	44
3.2.1	OS/HW Controlled P-states	44
3.2.1.1	Enhanced Intel SpeedStep® Technology	44
3.2.1.2	Intel® Speed Shift Technology	44
3.2.2	Low-Power Idle States	44
3.2.3	Requesting Low-Power Idle States	45
3.2.4	Processor IA Core C-State Rules	45
3.2.5	Package C-States	46
3.2.6	Package C-States and Display Resolutions	49
3.3	Processor Graphics Power Management	50
3.3.1	Memory Power Savings Technologies	50
3.3.1.1	Intel® Rapid Memory Power Management (Intel® RMPM)	50
3.3.2	Display Power Savings Technologies	50
3.3.2.1	Intel® Seamless Display Refresh Rate Switching Technology (Intel® SDRRS Technology) with eDP* Port	50
3.3.2.2	Intel® Automatic Display Brightness	51
3.3.2.3	Smooth Brightness	51
3.3.2.4	Intel® Display Power Saving Technology (Intel® DPST) 6.3	51
3.3.2.5	Panel Self-Refresh 2 (PSR 2)	51
3.3.2.6	Low-Power Single Pipe (LPSP)	52
3.3.2.7	Intel® Smart 2D Display Technology (Intel® S2DDT)	52
3.3.3	Processor Graphics Core Power Savings Technologies	52
3.3.3.1	Intel® Graphics Dynamic Frequency	52
3.3.3.2	Intel® Graphics Render Standby Technology (Intel® GRST)	52
3.3.3.3	Dynamic FPS (DFPS)	53
3.4	System Agent Enhanced Intel SpeedStep® Technology	53
3.5	Voltage Optimization	53
3.6	ROP (Rest Of Platform) PMIC	53
4	Thermal Management	54
4.1	Y/U-Processor Line Thermal and Power Specifications	54
4.2	Processor Thermal Management	56
4.2.1	Thermal Considerations	56
4.2.1.1	Intel® Turbo Boost Technology 2.0 availability may vary between the different SKUs. Package Power Control	57
4.2.1.2	Platform Power Control	58
4.2.1.3	Turbo Time Parameter (Tau)	59
4.2.2	Configurable TDP (cTDP) and Low-Power Mode	59
4.2.2.1	Configurable TDP	59
4.2.2.2	Low-Power Mode	60
4.2.3	Thermal Management Features	61
4.2.3.1	Adaptive Thermal Monitor	61
4.2.3.2	Digital Thermal Sensor	63



4.2.3.3	PROCHOT# Signal.....	64
4.2.3.4	PROCHOT Input Only.....	64
4.2.3.5	PROCHOT Output Only.....	65
4.2.3.6	Bi-Directional PROCHOT#	65
4.2.3.7	PROCHOT Demotion Algorithm.....	65
4.2.3.8	Voltage Regulator Protection using PROCHOT#	66
4.2.3.9	Thermal Solution Design and PROCHOT# Behavior	66
4.2.3.10	Low-Power States and PROCHOT# Behavior	66
4.2.3.11	THRMTRIP# Signal.....	67
4.2.3.12	Critical Temperature Detection	67
4.2.3.13	On-Demand Mode.....	67
4.2.3.14	I/O Emulation-Based On-Demand Mode	67
4.2.4	Intel® Memory Thermal Management.....	67
5	Memory	68
5.1	System Memory Interface	68
5.1.1	Processor SKU Support Matrix.....	68
5.1.1.1	LPDDR4/x Supported Memory Modules and Devices.....	71
5.1.2	System Memory Timing Support.....	71
5.1.3	System Memory Controller Organization Modes	72
5.1.4	System Memory Frequency.....	73
5.1.5	Technology Enhancements of Intel® Fast Memory Access (Intel® FMA).....	74
5.1.6	Data Scrambling	74
5.1.7	Data Swapping	74
5.1.8	DDR I/O Interleaving	75
5.1.9	Data Swapping	76
5.1.10	DRAM Clock Generation.....	76
5.1.11	DRAM Reference Voltage Generation	76
5.1.12	Data Swizzling	76
5.2	Integrated Memory Controller (IMC) Power Management	76
5.2.1	Disabling Unused System Memory Outputs	77
5.2.2	DRAM Power Management and Initialization	77
5.2.2.1	Initialization Role of CKE.....	78
5.2.2.2	Conditional Self-Refresh.....	78
5.2.2.3	Dynamic Power-Down.....	79
5.2.2.4	DRAM I/O Power Management.....	79
5.2.3	DDR Electrical Power Gating	79
5.2.4	Power Training.....	79
6	USB-C* Sub System	80
6.1	General Characteristics	80
6.2	USB3.x Supported Features.....	80
6.3	TCSS USB Blocks	81
6.3.1	USB Controllers.....	81
6.3.2	PHY.....	81
6.3.3	Integrated Thunderbolt™	82
6.4	Power states.....	83
7	Thunderbolt™	84
7.1	Host Router Implementation Capabilities	84
8	Graphics	86
8.1	Processor Graphics	86
8.1.1	Media Support (Intel® QuickSync and Clear Video Technology HD).....	86
8.1.1.1	Hardware Accelerated Video Decode.....	86
8.1.1.2	Hardware Accelerated Video Encode	87
8.1.1.3	Hardware Accelerated Video Processing	88
8.1.1.4	Hardware Accelerated Transcoding	88



8.2	Platform Graphics Hardware Feature	88
8.2.1	Hybrid Graphics.....	88
9	Display	90
9.1	Display Technologies Support.....	90
9.2	Display Configuration	90
9.3	Display Features.....	91
9.3.1	General Capabilities	91
9.3.2	Multiple Display Configurations	92
9.3.3	High-bandwidth Digital Content Protection (HDCP)	92
9.3.4	DisplayPort*	92
9.3.4.1	Multi-Stream Transport (MST).....	93
9.3.5	High-Definition Multimedia Interface (HDMI*).....	94
9.3.6	Digital Video Interface (DVI)	96
9.3.7	embedded DisplayPort* (eDP*)	96
9.3.8	Integrated Audio	97
10	Camera/MIPI	98
10.1	Camera Pipe Support	98
10.2	MIPI* CSI-2 Camera Interconnect.....	98
10.2.1	Camera Control Logic.....	98
10.2.2	Camera Modules.....	98
10.2.3	CSI-2 Lane Configuration.....	99
11	Signal Description	100
11.1	System Memory Interface.....	100
11.1.1	DDR4 Memory Interface	100
11.1.2	LPDDR4 Memory Interface	102
11.2	Reset and Miscellaneous Signals	103
11.3	Display Interfaces.....	104
11.3.1	Embedded DisplayPort* (eDP*) Signals	104
11.3.2	Digital Display Interface (DDI) Signals.....	104
11.4	USB Type-C Signals	105
11.5	MIPI* CSI-2 Interface Signals	105
11.6	Testability Signals.....	106
11.7	Error and Thermal Protection Signals.....	107
11.8	Power Sequencing Signals.....	107
11.9	Processor Power Rails	108
11.10	Ground, Reserved and Non-Critical to Function (NCTF) Signals.....	109
11.11	Processor Internal Pull-Up / Pull-Down Terminations.....	110
12	Electrical Specifications	111
12.1	Processor Power Rails	111
12.1.1	Power and Ground Pins.....	111
12.1.2	Integrated Voltage Regulator.....	111
12.1.3	V _{CC} Voltage Identification (VID).....	112
12.2	DC Specifications	112
12.2.1	Processor Power Rails DC Specifications	113
12.2.1.1	V _{CC1N} DC Specifications	113
12.2.1.2	V _{CC1P8A} DC Specifications.....	114
12.2.1.3	V _{CCIN_AUX} DC Specifications	114
12.2.1.4	V _{DDQ} DC Specifications	115
12.2.1.5	V _{CCST} DC Specifications	116
12.2.1.6	V _{CCPLL} DC Specifications	117
12.2.2	Processor Interfaces DC Specifications.....	118
12.2.2.1	DDR4 DC Specifications.....	118
12.2.2.2	LPDDR4/x DC Specifications.....	119



12.2.2.3	Digital Display Interface (DDI) DC Specifications	120
12.2.2.4	embedded DisplayPort* (eDP*) DC Specification	121
12.2.2.5	MIPI* CSI-2 D-Phy Receiver DC Specifications	121
12.2.2.6	CMOS DC Specifications.....	121
12.2.2.7	GTL and OD DC Specification.....	122
12.2.2.8	PECI DC Characteristics	122
12.3	Test Access Port (TAP) Connection.....	123
13	Package Mechanical Specifications	124
13.1	Package Mechanical Attributes	124
13.2	Package Loading and Die Pressure Specifications	124
13.2.1	Package Loading Specifications	125
13.2.2	Die Pressure Specifications	125
13.3	Package Storage Specifications	126
14	CPU And Device IDs	127
14.1	CPUID	127
14.2	PCI Configuration Header	128
14.3	Device IDs	128

Figures


1-1	U-Processor Line and Y-Processor Line Platforms	12
2-1	Example for PECI Host-Clients Connection	21
2-2	Example for PECI EC Connection	21
2-3	Device to Domain Mapping Structures.....	24
2-4	Processor Cache Hierarchy	33
2-5	Processor Camera System	40
3-1	Processor Power States.....	42
3-2	Idle Power Management Breakdown of the Processor IA Cores	44
3-3	Package C-State Entry and Exit.....	47
4-1	Package Power Control.....	58
4-2	PROCHOT Demotion Signal Description	66
5-1	Intel® Flex Memory Technology Operations	73
5-2	Interleave (IL) and Non-Interleave (NIL) Modes Mapping.....	76
6-1	USB-C* Sub-system Block Diagram	82
7-1	High Level Block Diagram.....	85
9-1	Processor Display Architecture	91
9-2	DisplayPort* Overview	93
9-3	HDMI* Overview	95
12-1	Input Device Hysteresis	123

Tables

1-1	Processor Lines	12
1-2	Terminology	16
1-3	Special marks	19
3-1	System States	43
3-2	Integrated Memory Controller (IMC) States	43
3-3	G, S, and C Interface State Combinations.....	43
3-4	Core C-states	46
3-5	Package C-States	47
3-6	Deepest Package C-State Available.....	49
3-7	Deepest Package C-State Available.....	50
4-1	TDP Specifications (U/Y-Processor Line)	55



4-2	Package Turbo Specifications	55
4-3	Junction Temperature Specifications.....	56
4-4	Configurable TDP Modes	59
5-1	DDR Support Matrix Table.....	68
5-2	DDR technology Support Matrix	68
5-5	SA Speed Enhanced Speed Steps (SA-GV) and Gear Mode Frequencies	69
5-3	DDR Max Capacity per System	69
5-4	LPDDR4/x Sub-Channels Population Rules.....	69
5-6	Supported DDR4 Non-ECC SODIMM Module Configurations (U-Processor Line)	70
5-7	Supported DDR4 Memory Down Device Configurations (U-Processor Line)	70
5-8	Supported LPDDR4/x x32 DRAMs Configurations (Y/U-Processor Line)	71
5-9	Supported LPDDR4/x x64 DRAMs Configurations (U/Y-Processor Line)	71
5-10	DDR4 System Memory Timing Support	72
5-11	LPDDR4/x System Memory Timing Support	72
5-12	Interleave (IL) and Non-Interleave (NIL) Modes Pin Mapping	75
6-1	USB Specifications	81
6-2	USB-C* Supported Configuration	81
6-3	USB-C* Non-Supported Configuration	82
6-4	PCIe* via TBT Configuration.....	82
6-5	TCSS power states.....	83
8-1	Supported configuration by SKU	86
8-2	Hardware Accelerated Video Decoding	87
8-3	Hardware Accelerated Video Encode	87
8-4	Hybrid Graphics Hardware Configuration.....	89
9-1	Display Ports Availability and Link Rate for Y/U-Processor Lines.....	90
9-2	Display Resolutions and Link Bandwidth for Multi-Stream Transport Calculations	93
9-3	DisplayPort* Maximum Resolution.....	94
9-4	HDMI* Maximum Resolution.....	96
9-5	DVI Maximum Resolution Supported.....	96
9-6	Embedded DisplayPort Maximum Resolution	96
9-7	Processor Supported Audio Formats over HDMI and DisplayPort*	97
11-1	Signal Tables Terminology	100
11-2	DDR4 Memory Interface	100
11-3	LPDDR4 Memory Interface	102
11-4	Reset and Miscellaneous Signals	103
11-5	embedded DisplayPort* Signals	104
11-6	Display Interface Signals	104
11-7	USB Type-C Signals	105
11-8	MIPI* CSI-2 Interface Signals	105
11-9	Testability Signals.....	106
11-10	Error and Thermal Protection Signals.....	107
11-11	Power Sequencing Signals	107
11-12	Processor Power Rails Signals.....	108
11-13	Processor Pull-up Power Rails Signals	109
11-14	GND, RSVD, and NCTF Signals	110
11-15	Processor Internal Pull-Up / Pull-Down Terminations.....	110
12-1	Processor Vcc _{IN} Active and Idle Mode DC Voltage and Current Specifications	113
12-2	Processor Vcc _{1p8A} Supply DC Voltage and Current Specifications.....	114
12-3	Vcc _{IN_AUX} Supply DC Voltage and Current Specifications	114
12-4	Memory Controller (VDDQ) Supply DC Voltage and Current Specifications	115
12-5	Vcc Sustain (Vcc _{ST}) Supply DC Voltage and Current Specifications	116
12-6	Vcc Sustain Gated (Vcc _{STG}) Supply DC Voltage and Current Specifications.....	116
12-7	Processor PLL (Vcc _{PLL}) Supply DC Voltage and Current Specifications.....	117
12-8	Processor PLL_OC (Vcc _{PLL_OC}) Supply DC Voltage and Current Specifications	117
12-9	DDR4 Signal Group DC Specifications	118




12-10LPDDR4/x Signal Group DC Specifications.....	119
12-11DSI HS Transmitter DC Specifications	120
12-12DSI LP Transmitter DC Specifications	120
12-13Digital Display Interface Group DC Specifications (DP/HDMI).....	120
12-14embedded DisplayPort* (eDP*) Group DC Specifications.....	121
12-15CMOS Signal Group DC Specifications	121
12-16 GTL Signal Group and Open Drain Signal Group DC Specifications.....	122
12-17PECI DC Electrical Limits	122
13-1 Package Mechanical Attributes	124
13-2 Package Loading Specifications	125
14-1 CPUID Format.....	127
14-2 Component Identification	127
14-3 PCI Configuration Header	128
14-4 Host Device ID (DID0)	128
14-5 Other Device ID	128



Revision History

Revision Number	Description	Revision Date
001	Initial Release	August 2019
002	<ul style="list-style-type: none">Added PECI clarification note Section 2.1, "Platform Environmental Control Interface (PECI)"Added Technology explanation Section , ""Added a note to Section 4.2.3.2, "Digital Thermal Sensor"Added TT1/TT2 feature explanation Section 4.2.3.1.4, "TT2/TT1 (Thermal Throttling Point)"Added power states clarifications Section 6.4, "Power states"	November 2019
003	<ul style="list-style-type: none">Updated Table 4-1, "TDP Specifications (U/Y-Processor Line)"Updated the I_{LI} value in Table 12-16	April 2020
004	<ul style="list-style-type: none">Removed references to TSX-NIMoved Device IDs to a new section. Section 14.3, "Device IDs"Update table 5-5Update 12.1.3 content	May 2020
005	<ul style="list-style-type: none">Updated Table 5-5Updated Section 12.1.3 for PEACH PRT changes	June 2020



1 Introduction

The 10th Generation Intel® Core™ processor is a 64-bit, multi-core processor built on 10-nanometer process technology.

The U-Processor Line and Y-Processor Line processors are offered in a 1-Chip Platform that includes the Intel® 495 Series Chipset Family On-Package Platform Controller Hub die on the same package as the processor die. Refer the following figure. The following table describes the different processor lines:

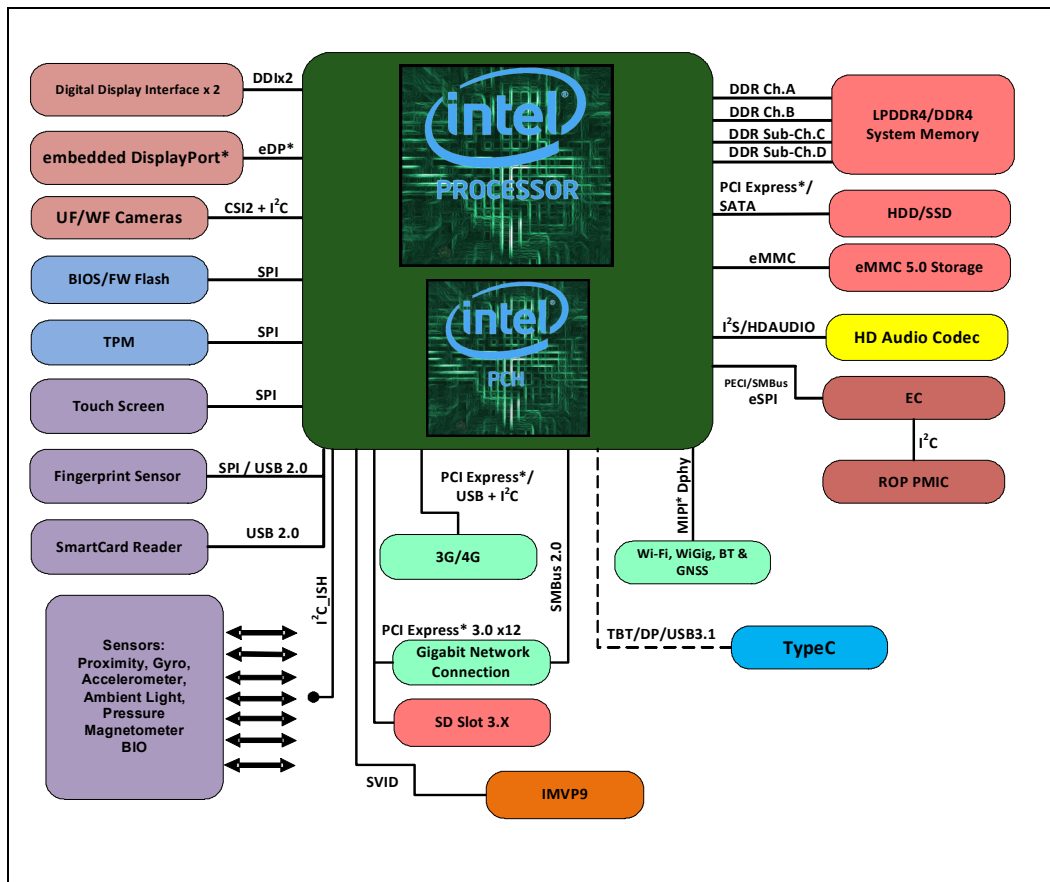
Table 1-1. Processor Lines

Processor Line ¹	Package	Base TDP	Processor IA Cores	EUs	VDBox	Platform Type
Y-Processor Line	BGA1377	9W	4	64/48/32	2/1	1-Chip
U-Processor Line	BGA1526	15W	4	64/48/32	2/1	
U-Processor Line	BGA1526	15W	2	32	1	

Notes:

1. Processor lines offering may change.
2. For additional TDP Configuration refer to Table 4-1, "TDP Specifications (U/Y-Processor Line)"
3. TDP workload does not reflect various I/O connectivity cases such as Thunderbolt™.

Figure 1-1. U-Processor Line and Y-Processor Line Platforms





This document covers all 10th Generation Intel® Core™ client segments processor lines (U and Y) for client segment. Not all processor interfaces and features are present in All Processor Lines. The presence of various interfaces and features will be indicated within the relevant sections and tables.

Throughout this document, the 10th Generation Intel® Core™ processor may be referred to simply as “processor” and the Intel® 495 Series Chipset Family On-Package Platform Controller Hub may be referred to simply as “PCH”.

1.1 Processor Volatility Statement

10th Generation Intel® Core™ processor families do not retain any end user data when powered down and/or when the processor is physically removed.

Note: Powered down refers to state which all processor power rails are off.

1.2 Package Support

The processor is available in the following packages:

- A 26.5 x 18.5 mm BGA package for Y-Processor Line
- A 50 x 25 mm BGA package for U-Processor Line

1.3 Supported Technologies

- Intel® Virtualization Technology (Intel® VT)
- Intel® Trusted Execution Technology (Intel® TXT)
- Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)
- PCLMULQDQ (Perform Carry-Less Multiplication Quad word) Instruction
- Intel® Secure Key
- Execute Disable Bit
- Intel® Boot Guard
- SMEP – Supervisor Mode Execution Protection
- SMAP – Supervisor Mode Access Protection
- Intel® Software Guard Extensions (Intel® SGX)
- SHA Extensions – Secure Hash Algorithm Extensions
- UMIP – User Mode Instruction Prevention
- RDPID – Read Processor ID
- Intel® Hyper-Threading Technology (Intel® HT Technology)
- Intel® Turbo Boost Technology 2.0
- Intel® Turbo Boost Max Technology 3.0
- Intel® SpeedStep Technology
- Intel® Speed Shift Technology
- Intel® Advanced Vector Extensions 2 (Intel® AVX2)



- Intel® 64 Architecture x2APIC
- PAIR – Power Aware Interrupt Routing
- Intel® GNA (GMM and Neural Network Accelerator)
- Intel® Image Processing Unit (Intel® IPU)
- Intel® Processor Trace
- PECI – Platform Environmental Control Interface

Note: The availability of the features may vary between processor SKUs.
Refer to [Chapter 2, “Technologies”](#) for more information.

1.3.1 API Support (Windows*)

- Direct3D* 2015, Direct3D* 12, Direct3D* 11.2, Direct3D* 11.1, Direct3D* 9, Direct3D* 10, Direct2D*
- OpenGL* 4.5
- OpenCL* 2.1, OpenCL 2.0, OpenCL 1.2

DirectX* extensions:

- PixelSync, InstantAccess, Conservative Rasterization, Render Target Reads, Floating-point De-norms, Shared Virtual memory, Floating Point atomics, MSAA sample-indexing, Fast Sampling (Coarse LOD), Quilted Textures, GPU Enqueue Kernels, GPU Signals processing unit. Other enhancements include color compression.

Gen 11 architecture delivers hardware acceleration of Direct X* 12 Render pipeline comprising of the following stages: Vertex Fetch, Vertex Shader, Hull Shader, Tessellation, Domain Shader, Geometry Shader, Rasterizer, Pixel Shader, Pixel Output.

1.4 Power Management Support

1.4.1 Processor Core Power Management

- Full support of ACPI C-states as implemented by the following processor C-states:
 - C0, C1, C1E, C6, C7, C8, C9, C10
- Enhanced Intel SpeedStep® Technology
- Intel® Speed Shift Technology

Refer to [Section 3.2, “Processor IA Core Power Management”](#) for more information.

1.4.2 System Power Management

- S0/S0ix, S3, S4, S5

Refer to [Chapter 3, “Power Management”](#) for more information.

1.4.3 Memory Controller Power Management

- Disabling Unused System Memory Outputs



- DRAM Power Management and Initialization
- Initialization Role of CKE
- Conditional Self-Refresh
- Dynamic Power Down
- DRAM I/O Power Management
- DDR Electrical Power Gating (EPG)
- Power training

Refer to [Section 5.2, “Integrated Memory Controller \(IMC\) Power Management”](#) for more information.

1.4.4 Processor Graphics Power Management

1.4.4.1 Memory Power Savings Technologies

- Intel® Rapid Memory Power Management (Intel® RMPM)
- Intel® Smart 2D Display Technology (Intel® S2DDT)

1.4.4.2 Display Power Savings Technologies

- Intel® (Seamless and Static) Display Refresh Rate Switching (DRRS) with eDP port
- Intel® Automatic Display Brightness
- Smooth Brightness
- Intel® Display Power Saving Technology (Intel® DPST 6)
- Panel Self-Refresh 2 (PSR 2)
- Low Power Single Pipe (LPSP)

1.4.4.3 Graphics Core Power Savings Technologies

- Intel® Graphics Dynamic Frequency
- Intel® Graphics Render Standby Technology (Intel® GRST)
- Dynamic FPS (Intel® DFPS)

Refer to [Section 3.3, “Processor Graphics Power Management”](#) for more information.

1.5 Thermal Management Support

- Digital Thermal Sensor
- Intel® Adaptive Thermal Monitor
- THERMTRIP# and PROCHOT# support
- On-Demand Mode
- Memory Open and Closed Loop Throttling
- Memory Thermal Throttling
- External Thermal Sensor (TS-on-DIMM and TS-on-Board)
- Render Thermal Throttling



- Fan speed control with DTS
- Intel® Turbo Boost Technology 2.0 Power Control

Refer to [Chapter 4, “Thermal Management”](#) for more information.

1.6 Processor Testability

An LTB on-board connector should be placed, to enable full debug capabilities. For the processor SKUs, a DCI (Direct Connect Interface) Tool is highly recommended to enable lower C-state debug.

1.7 Operating Systems Support

Processor Line	Windows* 10 64-bit	OS X	Linux* OS	Chrome* OS
Y-processor line	Yes	Yes	Yes	Yes
U-processor line	Yes	Yes	Yes	Yes

1.8 Terminology and Special Marks

Table 1-2. Terminology (Sheet 1 of 3)

Term	Description
4K	Ultra High Definition (UHD)
AES	Advanced Encryption Standard
AGC	Adaptive Gain Control
AVC	Advanced Video Coding
BLT	Block Level Transfer
BPP	Bits per pixel
CDR	Clock and Data Recovery
CTLE	Continuous Time Linear Equalizer
DDC	Digital Display Channel
DDI	Digital Display Interface for DP or HDMI/DVI
DSI	Display Serial Interface
DDR4	Fourth-Generation Double Data Rate SDRAM Memory Technology
DFE	Decision feedback equalizer
DMA	Direct Memory Access
DPPM	Dynamic Power Performance Management
DP*	DisplayPort*
DSC	Display Stream Compression
DSI	Display Serial Interface
DTS	Digital Thermal Sensor
ECC	Error Correction Code - used to fix DDR transactions errors
eDP*	embedded DisplayPort*
EU	Execution Unit in the Processor Graphics
FIVR	Fully Integrated Voltage Regulator



Table 1-2. Terminology (Sheet 2 of 3)

Term	Description
GSA	Graphics in System Agent
HDCP	High-bandwidth Digital Content Protection
HDMI*	High Definition Multimedia Interface
IMC	Integrated Memory Controller
Intel® 64 Technology	64-bit memory extensions to the IA-32 architecture
Intel® DPST	Intel® Display Power Saving Technology
Intel® PTT	Intel® Platform Trust Technology
Intel® TSX-NI	Intel® Transactional Synchronization Extensions
Intel® TXT	Intel® Trusted Execution Technology
Intel® VT	Intel® Virtualization Technology. Processor Virtualization, when used in conjunction with Virtual Machine Monitor software, enables multiple, robust independent software environments inside a single platform.
Intel® VT-d	Intel® Virtualization Technology (Intel® VT) for Directed I/O. Intel® VT-d is a hardware assist, under system software (Virtual Machine Manager or OS) control, for enabling I/O device Virtualization. Intel® VT-d also brings robust security by providing protection from errant DMAs by using DMA remapping, a key feature of Intel® VT-d.
ITH	Intel® Trace Hub
IOV	I/O Virtualization
IPU	Image Processing Unit
LFM	Low Frequency Mode. corresponding to the Enhanced Intel SpeedStep® Technology's lowest voltage/frequency pair.
LLC	Last Level Cache
LPDDR4/x	Low Power Double Data Rate SDRAM memory technology /x- additional power save.
LPM	Low Power Mode. The LPM Frequency is less than or equal to the LFM Frequency. The LPM TDP is lower than the LFM TDP as the LPM configuration limits the processor to single thread operation
LPSP	Low-Power Single Pipe
LSF	Lowest Supported Frequency. This frequency is the lowest frequency where manufacturing confirms logical functionality under the set of operating conditions.
MCP	Multi Chip Package - includes the processor and the PCH. In some SKUs it might have additional On-Package Cache.
MFM	Minimum Frequency Mode. MFM is the minimum ratio supported by the processor.
MLC	Mid-Level Cache
MPEG	Motion Picture Expert Group, international standard body JTC1/SC29/WG11 under ISO/IEC that has defined audio and video compression standards such as MPEG-1, MPEG-2, and MPEG-4, etc.
NCTF	Non-Critical to Function. NCTF locations are typically redundant ground or non-critical reserved balls/lands, so the loss of the solder joint continuity at end of life conditions will not affect the overall product functionality.
OPVR	On-Package Voltage Regulator
PCH	Platform Controller Hub. The chipset with centralized platform capabilities including the main I/O interfaces along with display connectivity, audio features, power management, manageability, security, and storage features. The PCH may also be referred as "chipset".
PECI	Platform Environment Control Interface
PL1, PL2, PL3	Power Limit 1, Power Limit 2, Power Limit 3
PMIC	Power Management Integrated Circuit
Processor	The 64-bit multi-core component (package)



Table 1-2. Terminology (Sheet 3 of 3)

Term	Description
Processor Core	The term "processor core" refers to Si die itself, which can contain multiple execution cores. Each execution core has an instruction cache, data cache, and 256-KB L2 cache. All execution cores share the LLC.
Processor Graphics	Intel Processor Graphics
PSR	Panel Self-Refresh
PSx	Power Save States (PS0, PS1, PS2, PS3, PS4)
Rank	A unit of DRAM corresponding to four to eight devices in parallel, ignoring ECC. These devices are usually, but not always, mounted on a single side of a SoDIMM.
SCI	System Control Interrupt. SCI is used in the ACPI protocol.
SDP	Scenario Design Power
SGX	Software Guard Extension
SHA	Secure Hash Algorithm
SSC	Spread Spectrum Clock
SSIC	SuperSpeed Inter-Chip
Storage Conditions	Refer Section 13.3, "Package Storage Specifications" .
STR	Suspend to RAM
TAC	Thermal Averaging Constant
TBT	Thunderbolt™ Interface
TCC	Thermal Control Circuit
TDP	Thermal Design Power
TTV TDP	Thermal Test Vehicle TDP
V _{CC}	Processor core power supply
V _{CCGT}	Processor Graphics Power Supply
V _{CCIO}	I/O Power Supply
V _{CCSA}	System Agent Power Supply
VLD	Variable Length Decoding
VPID	Virtual Processor ID
V _{SS}	Processor Ground
D0ix-states	USB controller power states ranging from D0i0 to D0i3, where D0i0 is fully powered on and D0i3 is primarily powered off. Controlled by SW.
S0ix-states	Processor residency idle standby power states.
USB-R	The type of storage redirection used from AMT 11.0 onward. In contrast to IDE-R, which presents remote floppy or CD drives as though they were integrated in the host machine, USB-R presents remote drives as though they were connected via a USB port.

Table 1-3. Special marks (Sheet 1 of 2)

Mark	Definition
[]	Brackets ([]) sometimes follow a ball, pin, registers or bit name. These brackets enclose a range of numbers, for example TCP[2:0]_TXRX_P[1:0] may refer to 4 USB-C* pins or EAX[7:0] may indicate a range that is 8 bits in length.
_N / # / B	A suffix of _N or # or B indicates an active low signal. for example CATERR# Note: _N does not refer to differential pair of signals such as CLK_P,CLK_N
0x000	Hexadecimal numbers are identified with an x in the number. All numbers are decimal (base 10) unless otherwise specified. Non-obvious binary numbers have the 'b' enclosed at the end of the number for example 0101b



Table 1-3. Special marks (Sheet 2 of 2)

Mark	Definition
	A vertical blue bar in the outside margin of a page indicates that a changes was made since the previous revision of this document.

§ §



2 Technologies

This chapter provides a high-level description of Intel technologies implemented in the processor.

The implementation of the features may vary between the processor SKUs.

Details on the different technologies of Intel processors and other relevant external notes are located at the Intel technology web site: <http://www.intel.com/technology/>

2.1 Platform Environmental Control Interface (PECI)

PECI is an Intel proprietary interface that provides a communication channel between Intel processors and external components like Super IO (SIO) and Embedded Controllers (EC) to provide processor temperature, Turbo, Configurable TDP, and memory throttling control mechanisms and many other services. PECI is used for platform thermal management and real time control and configuration of processor features and performance.

Notes: PECI over eSPI is supported.

PECI GetTemp command response time may take up to 800 us. EC that configured to PECI short timeout may not be able to get PECI response on time at Connected/Modern standby since S0ix exit time is ~800 us. As a result, it is recommended to do GetTemp retries.

2.1.1 PECI Bus Architecture

The PECI architecture is based on a wired OR bus that the clients (as processor PECI) can pull up (with strong drive).

The idle state on the bus is '0' (logical low) and near zero (Logical voltage level).

The following figures demonstrates PECI design and connectivity:

- PECI Host-Clients Connection: While the host/originator can be third party PECI host and one of the PECI client is a processor PECI device.
- PECI EC Connection

Figure 2-1. Example for PECE Host-Clients Connection

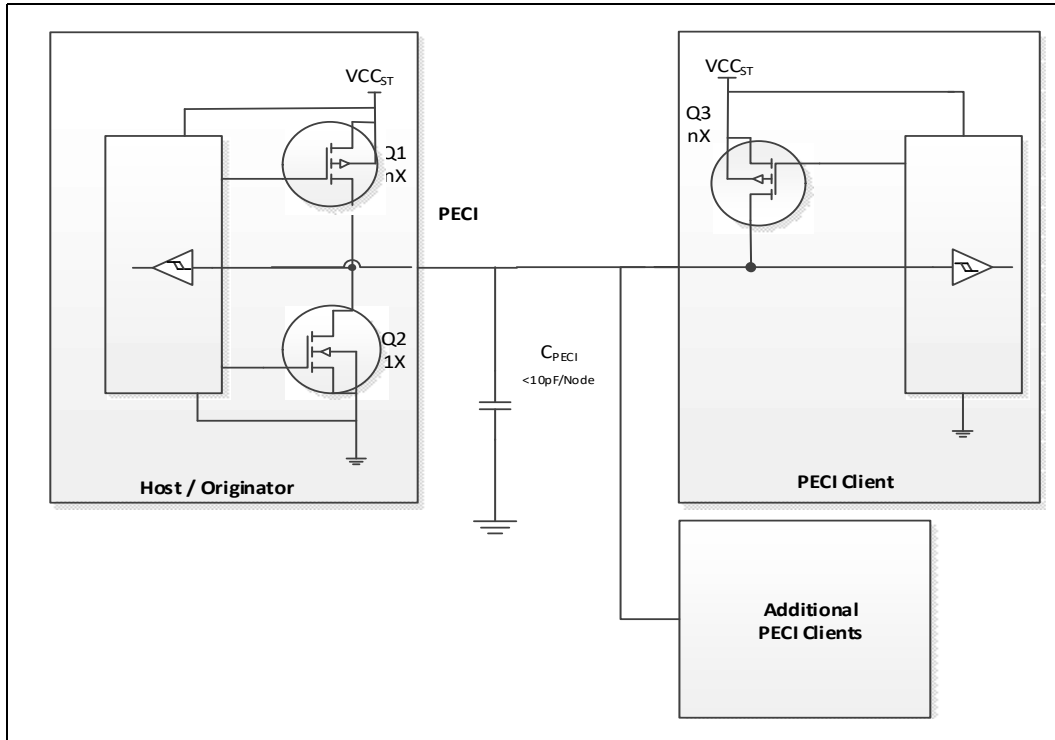
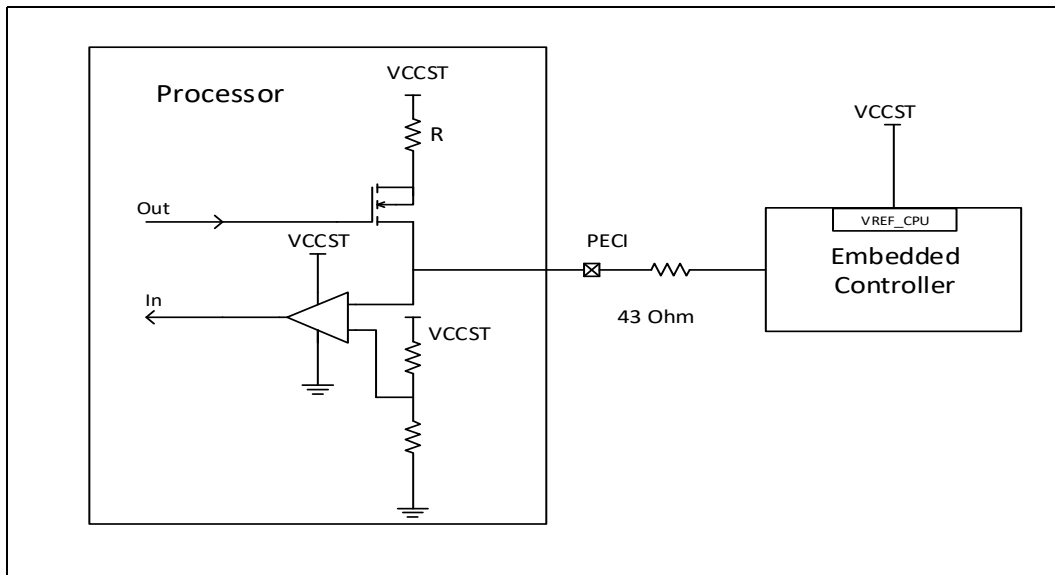


Figure 2-2. Example for PECE EC Connection





2.2 Intel® Virtualization Technology (Intel® VT)

Intel® Virtualization Technology (Intel® VT) makes a single system appear as multiple independent systems to software. This allows multiple, independent operating systems to run simultaneously on a single system. Intel® VT comprises technology components to support virtualization of platforms based on Intel architecture microprocessors and chipsets.

Intel® Virtualization Technology (Intel® VT) Intel® 64 and Intel® Architecture (Intel® VT-x) added hardware support in the processor to improve the virtualization performance and robustness. Intel® Virtualization Technology for Directed I/O (Intel® VT-d) extends Intel® VT-x by adding hardware assisted support to improve I/O device virtualization performance.

Intel® VT-x specifications and functional descriptions are included in the *Intel® 64 Architectures Software Developer's Manual, Volume 3*. Available at:

<http://www.intel.com/products/processor/manuals>

The Intel® VT-d specification and other VT documents can be referenced at:

<http://www.intel.com/content/www/us/en/virtualization/virtualization-technology/intel-virtualization-technology.html>

2.2.1 Intel® Virtualization Technology (Intel® VT) for Intel® 64 and Intel® Architecture (Intel® VT-X)

Intel® VT-x Objectives

Intel® VT-x provides hardware acceleration for virtualization of IA platforms. Virtual Machine Monitor (VMM) can use Intel® VT-x features to provide an improved reliable virtualized platform. By using Intel® VT-x, a VMM is:

- **Robust:** VMMs no longer need to use para-virtualization or binary translation. This means that VMMs will be able to run off-the-shelf operating systems and applications without any special steps.
- **Enhanced:** Intel® VT enables VMMs to run 64-bit guest operating systems on IA x86 processors.
- **More Reliable:** Due to the hardware support, VMMs can now be smaller, less complex, and more efficient. This improves reliability and availability and reduces the potential for software conflicts.
- **More Secure:** The use of hardware transitions in the VMM strengthens the isolation of VMs and further prevents corruption of one VM from affecting others on the same system.

Intel® VT-x Key Features

The processor supports the following added new Intel® VT-x features:

- Extended Page Table (EPT) Accessed and Dirty Bits
 - EPT A/D bits enabled VMMs to efficiently implement memory management and page classification algorithms to optimize VM memory operations, such as defragmentation, paging, live migration, and check-pointing. Without hardware support for EPT A/D bits, VMMs may need to emulate A/D bits by marking EPT paging-structures as not-present or read-only, and incur the overhead of EPT page-fault VM exits and associated software processing.
- EPTP (EPT pointer) switching

- EPTP switching is a specific VM function. EPTP switching allows guest software (in VMX non-root operation, supported by EPT) to request a different EPT paging-structure hierarchy. This is a feature by which software in VMX non-root operation can request a change of EPTP without a VM exit. Software will be able to choose among a set of potential EPTP values determined in advance by software in VMX root operation.
- Pause loop exiting
 - Support VMM schedulers seeking to determine when a virtual processor of a multiprocessor virtual machine is not performing useful work. This situation may occur when not all virtual processors of the virtual machine are currently scheduled and when the virtual processor in question is in a loop involving the PAUSE instruction. The new feature allows detection of such loops and is thus called PAUSE-loop exiting.

The processor IA core supports the following Intel® VT-x features:

- Extended Page Tables (EPT)
 - EPT is hardware assisted page table virtualization
 - It eliminates VM exits from guest OS to the VMM for shadow page-table maintenance
- Virtual Processor IDs (VPID)
 - Ability to assign a VM ID to tag processor IA core hardware structures (such as TLBs)
 - This avoids flushes on VM transitions to give a lower-cost VM transition time and an overall reduction in virtualization overhead.
- Guest Preemption Timer
 - Mechanism for a VMM to preempt the execution of a guest OS after an amount of time specified by the VMM. The VMM sets a timer value before entering a guest
 - The feature aids VMM developers in flexibility and Quality of Service (QoS) guarantees
- Descriptor-Table Exiting
 - Descriptor-table exiting allows a VMM to protect a guest OS from internal (malicious software based) attack by preventing relocation of key system data structures like IDT (interrupt descriptor table), GDT (global descriptor table), LDT (local descriptor table), and TSS (task segment selector).
 - A VMM using this feature can intercept (by a VM exit) attempts to relocate these data structures and prevent them from being tampered by malicious software.

2.2.2 Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d)

Intel® VT-d Objectives

The key Intel® VT-d objectives are domain-based isolation and hardware-based virtualization. A domain can be abstractly defined as an isolated environment in a platform to which a subset of host physical memory is allocated. Intel® VT-d provides accelerated I/O performance for a virtualized platform and provides software with the following capabilities:


- **I/O Device Assignment and Security:** for flexibly assigning I/O devices to VMs and extending the protection and isolation properties of VMs for I/O operations.
- **DMA Remapping:** for supporting independent address translations for Direct Memory Accesses (DMA) from devices.



- **Interrupt Remapping:** for supporting isolation and routing of interrupts from devices and external interrupt controllers to appropriate VMs.
- **Reliability:** for recording and reporting to system software DMA and interrupt errors that may otherwise corrupt memory or impact VM isolation.

Intel® VT-d accomplishes address translation by associating transaction from a given I/O device to a translation table associated with the Guest to which the device is assigned. It does this by means of the data structure in the following illustration. This table creates an association between the device's PCI Express* Bus/Device/Function (B/D/F) number and the base address of a translation table. This data structure is populated by a VMM to map devices to translation tables in accordance with the device assignment restrictions above, and to include a multi-level translation table (VT-d Table) that contains Guest specific address translations.

Figure 2-3. Device to Domain Mapping Structures



Intel® VT-d functionality, often referred to as an Intel® VT-d Engine, has typically been implemented at or near a PCI Express* host bridge component of a computer system. This might be in a chipset component or in the PCI Express* functionality of a



processor with integrated I/O. When one such VT-d engine receives a PCI Express* transaction from a PCI Express* bus, it uses the B/D/F number associated with the transaction to search for an Intel® VT-d translation table. In doing so, it uses the B/D/F number to traverse the data structure shown in the above figure. If it finds a valid Intel® VT-d table in this data structure, it uses that table to translate the address provided on the PCI Express bus. If it does not find a valid translation table for a given translation, this results in an Intel® VT-d fault. If Intel® VT-d translation is required, the Intel® VT-d engine performs an N-level table walk.

For more information, refer to Intel® Virtualization Technology for Directed I/O Architecture Specification <http://www.intel.com/content/dam/www/public/us/en/documents/product-specifications/vt-directed-io-spec.pdf>

Intel® VT-d Key Features

The processor supports the following Intel® VT-d features:

- Memory controller and processor graphics comply with the Intel® VT-d 2.1 Specification.
- Two Intel® VT-d DMA remap engines.
 - iGFX DMA remap engine
 - Default DMA remap engine (covers all devices except iGFX)
- Support for root entry, context entry, and default context
- 39-bit guest physical address and host physical address widths
- Support for 4 K page sizes only
- Support for register-based fault recording only (for single entry only) and support for MSI interrupts for faults
- Support for both leaf and non-leaf caching
- Support for boot protection of default page table
- Support for non-caching of invalid page table entries
- Support for hardware based flushing of translated but pending writes and pending reads, on IOTLB invalidation
- Support for Global, Domain specific and Page specific IOTLB invalidation
- MSI cycles (MemWr to address FEEx_xxxxh) not translated Interrupt Remapping is supported
- Queued invalidation is supported
- Intel® VT-d translation bypass address range is supported (Pass Through)

The processor supports the following added new Intel® VT-d features:

- 4-level Intel® VT-d Page walk – both default Intel® VT-d engine as well as the Processor Graphics VT-d engine are upgraded to support 4-level Intel® VT-d tables (adjusted guest address width of 48 bits)
- Intel® VT-d superpage – support of Intel® VT-d superpage (2 MB, 1 GB) for default Intel® VT-d engine (that covers all devices except IGD)
IGD Intel® VT-d engine does not support superpage and BIOS should disable superpage in default Intel® VT-d engine when iGfx is enabled.

Note: Intel® VT-d Technology may not be available on all SKUs.



2.2.3 Intel® APIC Virtualization Technology (Intel® APICv)

APIC virtualization is a collection of features that can be used to support the virtualization of interrupts and the Advanced Programmable Interrupt Controller (APIC).

When APIC virtualization is enabled, the processor emulates many accesses to the APIC, tracks the state of the virtual APIC, and delivers virtual interrupts — all in VMX non-root operation without a VM exit.

The following are the VM-execution controls relevant to APIC virtualization and virtual interrupts

- **Virtual-interrupt Delivery:** This control enables the evaluation and delivery of pending virtual interrupts. It also enables the emulation of writes to the APIC registers that control interrupt prioritization.
- **Use TPR Shadow:** This control enables emulation of accesses to the APIC's task-priority register (TPR) via CR8.
- **Virtualize APIC Accesses:** This control enables virtualization of memory-mapped accesses to the APIC by causing VM exits on accesses to a VMM-specified APIC-access page. Some of the other controls, if set, may cause some of these accesses to be emulated rather than causing VM exits.
- **Virtualize x2APIC Mode:** This control enables virtualization of accesses to the APIC.
- **APIC-register Virtualization:** This control allows memory-mapped of most APIC registers (as enabled) by satisfying them from the virtual-APIC page. It directs memory-mapped writes to the APIC-access page to the virtual-APIC page, following them by VM exits for VMM emulation.
- **Process Posted Interrupts:** This control allows software to post virtual interrupts in a data structure and send a notification to another logical processor; upon receipt of the notification, the target processor will process the posted interrupts by copying them into the virtual-APIC page.

Note: Intel® APIC Virtualization Technology may not be available on all SKUs.

Intel® APIC Virtualization specifications and functional descriptions are included in the *Intel® 64 Architectures Software Developer's Manual, Volume 3*. Available at:

<http://www.intel.com/products/processor/manuals>



2.3 Security Technologies

2.3.1 Intel® Trusted Execution Technology (Intel® TXT)

Intel® Trusted Execution Technology (Intel® TXT) defines platform-level enhancements that provide the building blocks for creating trusted platforms.

The Intel® TXT platform helps to provide the authenticity of the controlling environment such that those wishing to rely on the platform can make an appropriate trust decision. The Intel® TXT platform determines the identity of the controlling environment by accurately measuring and verifying the controlling software.

Another aspect of the trust decision is the ability of the platform to resist attempts to change the controlling environment. The Intel® TXT platform will resist attempts by software processes to change the controlling environment or bypass the bounds set by the controlling environment.

Intel® TXT is a set of extensions designed to provide a measured and controlled launch of system software that will then establish a protected environment for itself and any additional software that it may execute.

These extensions enhance two areas:

- The launching of the Measured Launched Environment (MLE).
- The protection of the MLE from potential corruption.

The enhanced platform provides these launch and control interfaces using Safer Mode Extensions (SMX).

The SMX interface includes the following functions:

- Measured/Verified launch of the MLE.
- Mechanisms to ensure the above measurement is protected and stored in a secure location.
- Protection mechanisms that allow the MLE to control attempts to modify itself.

The processor also offers additional enhancements to System Management Mode (SMM) architecture for enhanced security and performance. The processor provides new internal identifiers to:

- Enable a second SMM range
- Enable SMM code execution range checking
- Select whether SMM Save State is to be written to legacy SMRAM
- Determine if a thread is going to be delayed entering SMM
- Determine if a thread is blocked from entering SMM
- Targeted SMI, enable/disable threads from responding to SMIs, both VLWs and IPI

For the above features, BIOS should test the associated capability bit before attempting to access any of the above registers.

For more information, refer to the Intel® Trusted Execution Technology Measured Launched Environment Programming Guide at:

<http://www.intel.com/content/www/us/en/software-developers/intel-txt-software-development-guide.html>



Note: Intel® TXT Technology may not be available on all SKUs.

2.3.2 Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)

The processor supports Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI) that are a set of Single Instruction Multiple Data (SIMD) instructions that enable fast and secure data encryption and decryption based on the Advanced Encryption Standard (AES). Intel® AES-NI are valuable for a wide range of cryptographic applications, such as applications that perform bulk encryption/decryption, authentication, random number generation, and authenticated encryption. AES is broadly accepted as the standard for both government and industry applications, and is widely deployed in various protocols.

Intel® AES-NI consists of six Intel® SSE instructions. Four instructions, AESENC, AESENCLAST, AESDEC, and AESDELAST facilitate high performance AES encryption and decryption. The other two, AESIMC and AESKEYGENASSIST, support the AES key expansion procedure. Together, these instructions provide full hardware for supporting AES; offering security, high performance, and a great deal of flexibility.

This generation of the processor has increased the performance of the Intel® AES-NI significantly compared to previous products.

The Intel® AES-NI specifications and functional descriptions are included in the *Intel® 64 Architectures Software Developer's Manual, Volume 2*. Available at:

<http://www.intel.com/products/processor/manuals>

Note: Intel® AES-NI Technology may not be available on all SKUs.

2.3.3 Perform Carry-Less Multiplication Quad Word (PCLMULQDQ) Instruction

The processor supports the carry-less multiplication instruction, PCLMULQDQ. PCLMULQDQ is a Single Instruction Multiple Data (SIMD) instruction that computes the 128-bit carry-less multiplication of two 64-bit operands without generating and propagating carries. Carry-less multiplication is an essential processing component of several cryptographic systems and standards. Hence, accelerating carry-less multiplication can significantly contribute to achieving high speed secure computing and communication.

PCLMULQDQ specifications and functional descriptions are included in the *Intel® 64 Architectures Software Developer's Manual, Volume 2*. Available at:

<http://www.intel.com/products/processor/manuals>

2.3.4 Intel® Secure Key

The processor supports Intel® Secure Key (formerly known as Digital Random Number Generator (DRNG), a software visible random number generation mechanism supported by a high quality entropy source. This capability is available to programmers through the RDRAND instruction. The resultant random number generation capability is designed to comply with existing industry standards in this regard (ANSI X9.82 and NIST SP 800-90).



Some possible usages of the RDRAND instruction include cryptographic key generation as used in a variety of applications, including communication, digital signatures, secure storage, etc.

RDRAND specifications and functional descriptions are included in the *Intel® 64 Architectures Software Developer's Manual, Volume 2*. Available at:

<http://www.intel.com/products/processor/manuals>

2.3.5 Execute Disable Bit

The Execute Disable Bit allows memory to be marked as non-executable when combined with a supporting operating system. If code attempts to run in non-executable memory, the processor raises an error to the operating system. This feature can prevent some classes of viruses or worms that exploit buffer overrun vulnerabilities and can, thus, help improve the overall security of the system.

2.3.6 Boot Guard Technology

Boot Guard technology is a part of boot integrity protection technology. Boot Guard can help protect the platform boot integrity by preventing execution of unauthorized boot blocks. With Boot Guard, platform manufacturers can create boot policies such that invocation of an unauthorized (or untrusted) boot block will trigger the platform protection per the manufacturer's defined policy.

With verification based in the hardware, Boot Guard extends the trust boundary of the platform boot process down to the hardware level.

Boot Guard accomplishes this by:

- Providing of hardware-based Static Root of Trust for Measurement (S-RTM) and the Root of Trust for Verification (RTV) using Intel architectural components.
- Providing of architectural definition for platform manufacturer Boot Policy.
- Enforcing of manufacture provided Boot Policy using Intel architectural components.

Benefits of this protection is that Boot Guard can help maintain platform integrity by preventing re-purposing of the manufacturer's hardware to run an unauthorized software stack.

Note: Boot Guard availability may vary between the different SKUs.

2.3.7 Intel® Supervisor Mode Execution Protection (SMEP)

Intel® Supervisor Mode Execution Protection (SMEP) is a mechanism that provides the next level of system protection by blocking malicious software attacks from user mode code when the system is running in the highest privilege level. This technology helps to protect from virus attacks and unwanted code from harming the system.



2.3.8 Intel® Supervisor Mode Access Protection (SMAP)

Intel® Supervisor Mode Access Protection (SMAP) is a mechanism that provides next level of system protection by blocking a malicious user from tricking the operating system into branching off user data. This technology shuts down very popular attack vectors against operating systems.

2.3.9 Intel® Software Guard Extensions (Intel® SGX)

Software Guard Extensions (SGX) is a processor enhancement designed to help protect application integrity and confidentiality of secrets and withstands software and certain hardware attacks.

Software Guard Extensions (SGX) architecture provides the capability to create isolated execution environments named Enclaves that operate from a protected region of memory.

Enclave code can be accessed using new special ISA commands that jump into per Enclave predefined addresses. Data within an Enclave can only be accessed from that same Enclave code.

The latter security statements hold under all privilege levels including supervisor mode (ring-0), System Management Mode (SMM) and other Enclaves.

Intel® SGX features a memory encryption engine that both encrypt Enclave memory as well as protect it from corruption and replay attacks.

Intel® SGX benefits over alternative Trusted Execution Environments (TEEs) are:

- Enclaves are written using C/C++ using industry standard build tools.
- High processing power as they run on the processor.
- Large amount of memory are available as well as non-volatile storage (such as disk drives).
- Simple to maintain and debug using standard IDEs (Integrated Development Environment)
- Scalable to a larger number of applications and vendors running concurrently
- Dynamic memory allocation:
 - Heap and thread-pool management
 - On-demand stack growth
 - Dynamic module/library loading
 - Concurrency management in applications such as garbage collectors
 - Write-protection of EPC pages (Enclave Page Cache - Enclave protected memory) after initial relocation
 - On-demand creation of code pages (JIT, encrypted code modules)
- Allow Launch Enclaves other than the one currently provided by Intel
- Maximum protected memory size has increased to 256 MB.
 - Supports 64, 128 and 256 MB protected memory sizes.



- **VMM Over-subscription.** The VMM over-subscription mechanism allows a VMM to make more resources available to virtual machines than what is actually available on the platform. The initial Intel® SGX architecture was optimized for EPC partitioning/ballooning model for VMMs, where a VMM assigns a static EPC partition to each SGX guest OS without over-subscription and guests are free to manage (that is, oversubscribe) their own EPC partitions. The Intel® SGX EPC Over-subscription Extensions architecture provides a set of new instructions allowing VMMs to efficiently oversubscribe EPC memory for its guest operating systems.

For more information, refer to the Intel® SGX website at:

<https://software.intel.com/en-us/sgx>

2.3.10 Intel® Secure Hash Algorithm Extensions (Intel® SHA Extensions)

The Secure Hash Algorithm (SHA) is one of the most commonly employed cryptographic algorithms. Primary usages of SHA include data integrity, message authentication, digital signatures, and data de-duplication. As the pervasive use of security solutions continues to grow, SHA can be seen in more applications now than ever. The Intel® SHA Extensions are designed to improve the performance of these compute-intensive algorithms on Intel® architecture-based processors.

The Intel® SHA Extensions are a family of seven instructions based on the Intel® Streaming SIMD Extensions (Intel® SSE) that are used together to accelerate the performance of processing SHA-1 and SHA-256 on Intel architecture-based processors. Given the growing importance of SHA in our everyday computing devices, the new instructions are designed to provide a needed boost of performance to hashing a single buffer of data. The performance benefits will not only help improve responsiveness and lower power consumption for a given application, they may enable developers to adopt SHA in new applications to protect data while delivering to their user experience goals. The instructions are defined in a way that simplifies their mapping into the algorithm processing flow of most software libraries, thus enabling easier development.

More information on Intel® SHA can be found at:

<http://software.intel.com/en-us/articles/intel-sha-extensions>

2.3.11 User Mode Instruction Prevention (UMIP)

User Mode Instruction Prevention (UMIP) provides additional hardening capability to the OS kernel by allowing certain instructions to execute only in supervisor mode (Ring 0).

If the OS opt-in to use UMIP, the following instructions are enforced to run in supervisor mode:

- **SGDT:** Store the GDTR register value
- **SIDT:** Store the IDTR register value
- **SLDT:** Store the LDTR register value
- **SMSW:** Store Machine Status Word
- **STR:** Store the TR register value



An attempt at such execution in user mode causes a general protection exception (#GP).

2.3.12 Read Processor ID (RDPID)

A companion instruction that returns the current logical processor's ID and provides a faster alternative to using the RDTSCP instruction.

2.4 Power and Performance Technologies

2.4.1 Intel[®] Smart Cache Technology

The Intel[®] Smart Cache Technology is a shared Last Level Cache (LLC).

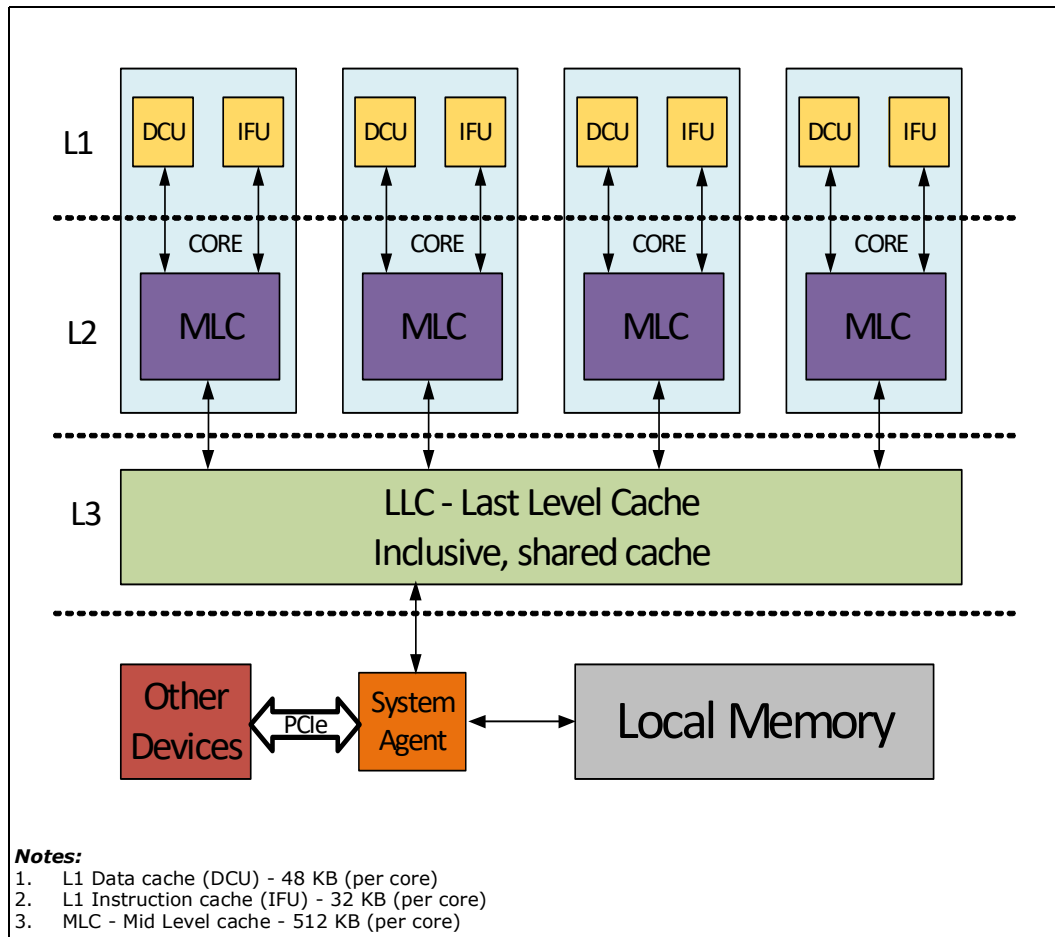
- The LLC may also be referred to as a third level cache.
- The LLC is shared between all IA cores as well as the Processor Graphics.
- The first and second level caches are not shared between physical cores and each physical core has a separate set of caches.
- The size of the LLC is SKU specific with a maximum of 2 MB per physical core and is a 16 way associative cache.

2.4.2 IA Core Level 1 and Level 2 Caches

The first-level cache is divided into a data cache and an instruction cache. The processor first level cache size is 48KB for data and 32KB for instructions. The first level cache is an eight way associative cache.

The second level cache holds both data and instructions. It is also referred to as mid-level cache or MLC. The processor second level cache size is 512 KB and is an eight way associative cache.

Figure 2-4. Processor Cache Hierarchy



2.4.3 Intel® Turbo Boost Max Technology 3.0

The Intel® Turbo Boost Max Technology 3.0 (ITBMT 3.0) grants a different maximum Turbo frequency for individual processor cores.

To enable ITBMT 3.0 the processor exposes individual core capabilities; including diverse maximum turbo frequencies.

An operating system that allows for varied per core frequency capability can then maximize power savings and performance usage by assigning tasks to the faster cores, especially on low core count workloads.

Processors enabled with these capabilities can also allow software (most commonly a driver) to override the maximum per-core Turbo frequency limit and notify the operating system via an interrupt mechanism.

For more information on the Intel® Turbo Boost Max 3.0 Technology, refer

<http://www.intel.com/content/www/us/en/architecture-and-technology/turbo-boost/turbo-boost-max-technology.html>



Note: Intel® Turbo Boost Max 3.0 Technology may not be available on all SKUs.

2.4.4 Power Aware Interrupt Routing (PAIR)

The processor includes enhanced power-performance technology that routes interrupts to threads or processor IA cores based on their sleep states. As an example, for energy savings, it routes the interrupt to the active processor IA cores without waking the deep idle processor IA cores. For performance, it routes the interrupt to the idle (C1) processor IA cores without interrupting the already heavily loaded processor IA cores. This enhancement is mostly beneficial for high-interrupt scenarios like Gigabit LAN, WLAN peripherals, and so on.

2.4.5 Intel® Hyper-Threading Technology (Intel® HT Technology)

The processor supports Intel® Hyper-Threading Technology (Intel® HT Technology) that allows an execution processor IA core to function as two logical processors. While some execution resources such as caches, execution units, and buses are shared, each logical processor has its own architectural state with its own set of general-purpose registers and control registers. This feature should be enabled using the BIOS and requires operating system support.

Intel recommends enabling Intel® Hyper-Threading Technology with Microsoft* Windows* 7 or newer and disabling Intel® Hyper-Threading Technology using the BIOS for all previous versions of Windows* operating systems. For more information on Intel® Hyper-Threading Technology, refer <http://www.intel.com/technology/platform-technology/hyper-threading/>

Note: Intel® HT Technology may not be available on all SKUs.

2.4.6 Intel® Turbo Boost Technology 2.0

The Intel® Turbo Boost Technology 2.0 allows the processor IA core / processor graphics core to opportunistically and automatically run faster than the processor IA core base frequency / processor graphics base frequency if it is operating below power, temperature, and current limits. The Intel® Turbo Boost Technology 2.0 feature is designed to increase performance of both multi-threaded and single-threaded workloads.

Compared with previous generation products, Intel® Turbo Boost Technology 2.0 will increase the ratio of application power towards TDP and also allows to increase power above TDP as high as PL2 for short periods of time. Thus, thermal solutions and platform cooling that are designed to less than thermal design guidance might experience thermal and performance issues since more applications will tend to run at the maximum power limit for significant periods of time.

Note: Intel® Turbo Boost Technology 2.0 may not be available on all SKUs.

2.4.6.1 Intel® Turbo Boost Technology 2.0 Power Monitoring

When operating in turbo mode, the processor monitors its own power and adjusts the processor and graphics frequencies to maintain the average power within limits over a thermally significant time period. The processor estimates the package power for all

components on package. In the event that a workload causes the temperature to exceed program temperature limits, the processor will protect itself using the Adaptive Thermal Monitor.

2.4.6.2 Intel® Turbo Boost Technology 2.0 Power Control

Illustration of Intel® Turbo Boost Technology 2.0 power control is shown in the following sections and figures. Multiple controls operate simultaneously allowing customization for multiple system thermal and power limitations. These controls allow for turbo optimizations within system constraints and are accessible using MMIO, and PECI interfaces.

2.4.6.3 Intel® Turbo Boost Technology 2.0 Frequency

To determine the highest performance frequency amongst active processor IA cores, the processor takes the following into consideration:

- The number of processor IA cores operating in the C0 state.
- The estimated processor IA core current consumption and I_{CCMax} settings.
- The estimated package prior and present power consumption and turbo power limits.
- The package temperature.

Any of these factors can affect the maximum frequency for a given workload. If the power, current, or thermal limit is reached, the processor will automatically reduce the frequency to stay within its TDP limit. Turbo processor frequencies are only active if the operating system is requesting the P0 state. For more information on P-states and C-states, refer [Chapter 3, “Power Management”](#).

2.4.7 Enhanced Intel SpeedStep® Technology

Enhanced Intel SpeedStep® Technology enables OS to control and select P-state. The following are the key features of Enhanced Intel SpeedStep® Technology:

- Multiple frequency and voltage points for optimal performance and power efficiency. These operating points are known as P-states.
- Frequency selection is software controlled by writing to processor internal identifiers. The voltage is optimized based on the selected frequency and the number of active processor IA cores.
 - Once the voltage is established, the PLL locks on to the target frequency.
 - All active processor IA cores share the same frequency and voltage. In a multi-core processor, the highest frequency P-state requested among all active IA cores is selected.
 - Software-requested transitions are accepted at any time. If a previous transition is in progress, the new transition is deferred until the previous transition is completed.
- The processor controls voltage ramp rates internally to ensure glitch-free transitions.

Note: Because there is low transition latency between P-states, a significant number of transitions per-second are possible.



2.4.8 Intel® Speed Shift Technology

Intel® Speed Shift Technology is an energy efficient method of frequency control by the hardware rather than relying on OS control. OS is aware of available hardware P-states and request a desired P-state or it can let the hardware determine the P-state. The OS request is based on its workload requirements and awareness of processor capabilities. Processor decision is based on the different system constraints for example: Workload demand, thermal limits while taking into consideration the minimum and maximum levels and activity window of performance requested by the Operating System.

2.4.9 Intel® Advanced Vector Extensions 2 (Intel® AVX2)

Intel® Advanced Vector Extensions 2.0 (Intel® AVX2) is the latest expansion of the Intel instruction set. Intel® AVX2 extends the Intel® Advanced Vector Extensions (Intel® AVX) with 256-bit integer instructions, floating-point fused multiply add (FMA) instructions, and gather operations. The 256-bit integer vectors benefit math, codec, image, and digital signal processing software. FMA improves performance in face detection, professional imaging, and high performance computing. Gather operations increase vectorization opportunities for many applications. In addition to the vector extensions, this generation of Intel processors adds new bit manipulation instructions useful in compression, encryption, and general purpose software.

For more information on Intel® AVX, refer <http://www.intel.com/software/avx>

Intel® Advanced Vector Extensions (Intel® AVX) are designed to achieve higher throughput to certain integer and floating point operation. Due to varying processor power characteristics, utilizing AVX instructions may cause a) parts to operate below the base frequency b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software and system configuration and user should consult the system manufacturer for more information.

Intel® Advanced Vector Extensions refers to Intel® AVX, Intel® AVX2 or Intel® AVX-512.

For more information on Intel® AVX, refer <https://software.intel.com/en-us/isa-extensions/intel-avx>.

Note: Intel® AVX and AVX2 Technologies may not be available on all SKUs.

2.4.10 Intel® 64 Architecture x2APIC

The x2APIC architecture extends the xAPIC architecture that provides key mechanisms for interrupt delivery. This extension is primarily intended to increase processor addressability.

Specifically, x2APIC:

- Retains all key elements of compatibility to the xAPIC architecture:
 - Delivery modes
 - Interrupt and processor priorities
 - Interrupt sources
 - Interrupt destination types
- Provides extensions to scale processor addressability for both the logical and physical destination modes



- Adds new features to enhance performance of interrupt delivery
- Reduces complexity of logical destination mode interrupt delivery on link based architectures

The key enhancements provided by the x2APIC architecture over xAPIC are the following:

- Support for two modes of operation to provide backward compatibility and extensibility for future platform innovations:
 - In xAPIC compatibility mode, APIC registers are accessed through memory mapped interface to a 4K-Byte page, identical to the xAPIC architecture.
- Increased range of processor addressability in x2APIC mode:
 - Physical xAPIC ID field increases from 8 bits to 32 bits, allowing for interrupt processor addressability up to 4G-1 processors in physical destination mode. A processor implementation of x2APIC architecture can support fewer than 32-bits in a software transparent fashion.
 - Logical xAPIC ID field increases from 8 bits to 32 bits. The 32-bit logical x2APIC ID is partitioned into two sub-fields – a 16-bit cluster ID and a 16-bit logical ID within the cluster. Consequently, $(2^{20} - 16)$ processors can be addressed in logical destination mode. Processor implementations can support fewer than 16 bits in the cluster ID sub-field and logical ID sub-field in a software agnostic fashion.
- The semantics for accessing APIC registers have been revised to simplify the programming of frequently-used APIC registers by system software. Specifically, the software semantics for using the Interrupt Command Register (ICR) and End Of Interrupt (EOI) registers have been modified to allow for more efficient delivery and dispatching of interrupts.
- The x2APIC extensions are made available to system software by enabling the local x2APIC unit in the “x2APIC” mode. To benefit from x2APIC capabilities, a new operating system and a new BIOS are both needed, with special support for x2APIC mode.
- The x2APIC architecture provides backward compatibility to the xAPIC architecture and forward extensible for future Intel platform innovations.

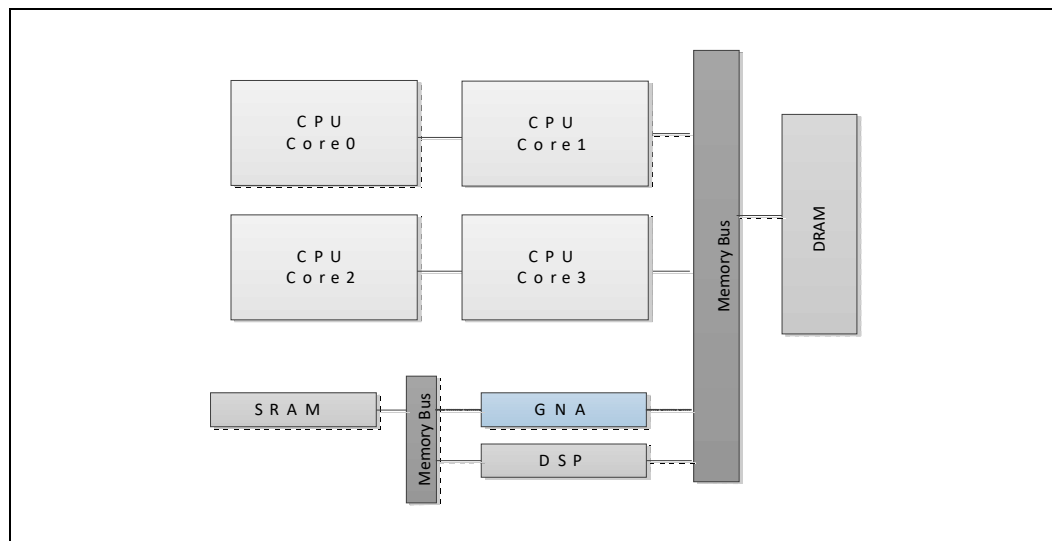
Note: Intel® x2APIC Technology may not be available on all SKUs.

For more information, refer the Intel® 64 Architecture x2APIC Specification at <http://www.intel.com/products/processor/manuals/>.

2.4.11 Intel® GNA (GMM and Neural Network Accelerator)

GNA stands for Gaussian Mixture Model and Neural Network Accelerator

The GNA is used to process speech recognition without user training sequence. The GNA is designed to unload the processor cores and the system memory with complex speech recognition tasks and improve the speech recognition accuracy. The GNA is designed to compute millions of Gaussian probability density functions per second without loading the processor cores while maintaining low power consumption.



2.4.12 Advanced Vector Extensions 512 Bit (Intel® AVX-512)

Intel® AVX support is widened to 512 bit SIMD operations. Programs can pack eight double precision and sixteen single precision floating numbers within the 512-bit vectors, as well as eight 64-bit and sixteen 32-bit integers. This enables processing of twice the number of data elements that Intel® AVX/AVX2 can process with a single instruction and four times the capabilities of Intel® SSE.

Intel® AVX-512 instructions are important because they open up higher performance capabilities for the most demanding computational tasks. Intel® AVX-512 instructions offer the highest degree of compiler support by including an unprecedented level of richness in the design of the instruction capabilities.

Intel® AVX-512 features include 32 vector registers each 512-bit wide and eight dedicated mask registers. Intel® AVX-512 is a flexible instruction set that includes support for broadcast, embedded masking to enable predication, embedded floating point rounding control, embedded floating-point fault suppression, scatter instructions, high speed math instructions, and compact representation of large displacement values.

Intel® AVX-512 offers a level of compatibility with Intel® AVX which is stronger than prior transitions to new widths for SIMD operations. Unlike Intel® SSE and Intel® AVX which cannot be mixed without performance penalties, the mixing of Intel® AVX and Intel® AVX-512 instructions is supported without penalty. Intel® AVX registers YMM0-YMM15 map into Intel® AVX-512 registers ZMM0-ZMM15 (in x86-64 mode), very much like Intel® SSE registers map into Intel® AVX registers. Therefore, in processors with Intel® AVX-512 support, Intel® AVX and Intel® AVX2 instructions operate on the lower 128 or 256 bits of the first 16 ZMM registers.

Intel® AVX-512 has multiple extensions that CPUID has been enhanced to expose.

- **AVX512F (Foundation)**: Expands most 32-bit and 64-bit based AVX instructions with EVEX coding scheme to support 512-bit registers, operation masks, parameter broadcasting, and embedded rounding and exception control



- **AVX512CD (Conflict Detection)**: Efficient conflict detection to allow more loops to be vectorized
- **AVX512BW (Byte and Word)**: Extends AVX-512 to cover 8-bit and 16-bit integer operations
- **AVX512DQ (Doubleword and Quadword)**: Extends AVX-512 to cover 32-bit and 64-bit integer operations
- **AVX512VL (Vector Length)**: Extends most AVX-512 operations to also operate on XMM (128-bit) and YMM (256-bit) registers
- **AVX512IFMA (Integer Fused Multiply-Add)**: Fused multiply-add of integers using 52-bit precision
- **AVX512VBMI (Vector Byte Manipulation Instructions)**: Adds vector byte permutation instructions which were not present in AVX-512BW
- **AVX512VBMI2 (Vector Byte Manipulation Instructions 2)**: Adds byte/word load, store and concatenation with shift
- **VPOPCNTDQ**: Count of bits set to 1
- **VPCLMULQDQ**: Carry-less multiplication of quadwords
- **AVX-512VNNI (Vector Neural Network Instructions)**: Vector instructions for deep learning
- **AVX512GFNI (Galois Field New Instructions)**: Vector instructions for calculating Galois Fields
- **AVX512VAES (Vector AES instructions)**: Vector instructions for AES coding
- **AVX512BITALG (Bit Algorithms)**: Byte/word bit manipulation instructions expanding VPOPCNTDQ

Note: Intel® AVX-512 may not be available on all SKUs.

2.4.13 Cache Line Write Back (CLWB)

Writes back to memory the cache line (if dirty) that contains the linear address specified with the memory operand from any level of the cache hierarchy in the cache coherence domain. The line may be retained in the cache hierarchy in non-modified state. Retaining the line in the cache hierarchy is a performance optimization (treated as a hint by hardware) to reduce the possibility of cache miss on a subsequent access. Hardware may choose to retain the line at any of the levels in the cache hierarchy, and in some cases, may invalidate the line from the cache hierarchy. The source operand is a byte memory location.

2.4.14 Ring Interconnect

The Ring is a high speed, wide interconnect that links the processor cores, processor graphics and the System Agent. The Ring shares frequency and voltage with the Last Level Cache (LLC). The Ring's frequency dynamically changes. Its frequency is relative to both processor cores and processor graphics frequencies.

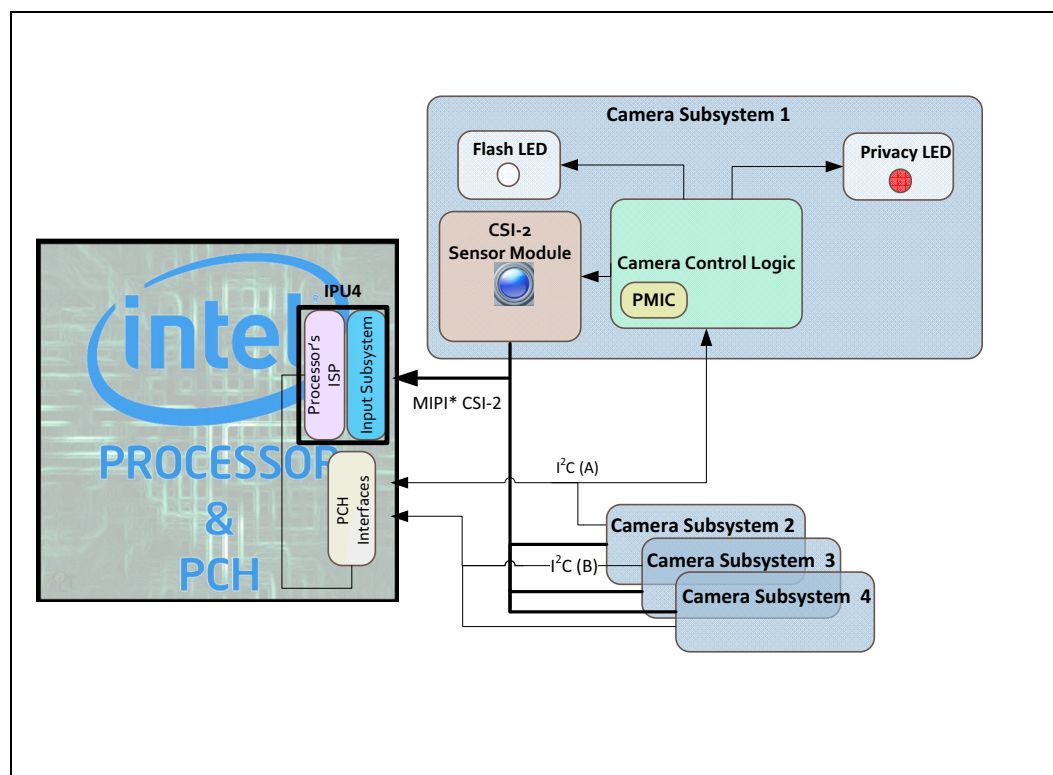
2.5 Intel® Image Processing Unit (Intel® IPU)

2.5.1 Platform Imaging Infrastructure

The platform imaging infrastructure is based on the following hardware components:

- **Camera SubSystem:** Located in the lid of the system and contains CMOS sensor, flash, LED, I/O interface (MIPI* CSI-2 and I²C*), Focus control and other components.
- **Camera I/O Controller:** The I/O controller is located in the processor and contains a MIPI-CSI2 Host controller. The host controller is a PCI device (independent of the IPU device). The CSI-2 HCI brings imaging data from an external image into the system and provides a command and control channel for the image using I²C.
- **Intel® IPU (Image Processing Unit):** The IPU processes raw images captured by Bayer sensors. The result images are used by still photography and video capture applications (JPEG, H.264, etc.).

Figure 2-5. Processor Camera System



2.5.2 Intel® Image Processing Unit (Intel® IPU)

The Intel® IPU is an embedded camera subsystem hardware component on the processor, it processes video and still images at high quality while consuming lower-power by leveraging a programmable VLIW (very-long-instruction-word) SIMD vector processor, a hardware fixed function pipe (accelerators), 3 scalar processors and more.



The mix of hardware accelerators and compute capabilities allow the flexibility and patch ability that are required for late changes and allows the processor to support future sensor technologies while maintaining both power and performance.

2.6 Debug Technologies

2.6.1 Intel® Processor Trace

Intel® Processor Trace (Intel® PT) is a tracing capability added to Intel® Architecture, for use in software debug and profiling. Intel® PT provides the capability for more precise software control flow and timing information, with limited impact to software execution. This provides enhanced ability to debug software crashes, hangs, or other anomalies, as well as responsiveness and short-duration performance issues.

Intel® VTune™ Amplifier for Systems and the Intel® System Debugger are part of Intel® System Studio 2015 (and newer) product, which includes updates for the new debug and trace features, including Intel® PT and Intel® Trace Hub.

Intel® System Studio is available for download at <https://software.intel.com/en-us/system-studio>.

An update to the Linux* performance utility, with support for Intel® PT, is available for download at https://github.com/virtuoso/linux-perf/tree/intel_pt. It requires rebuilding the kernel and the perf utility.



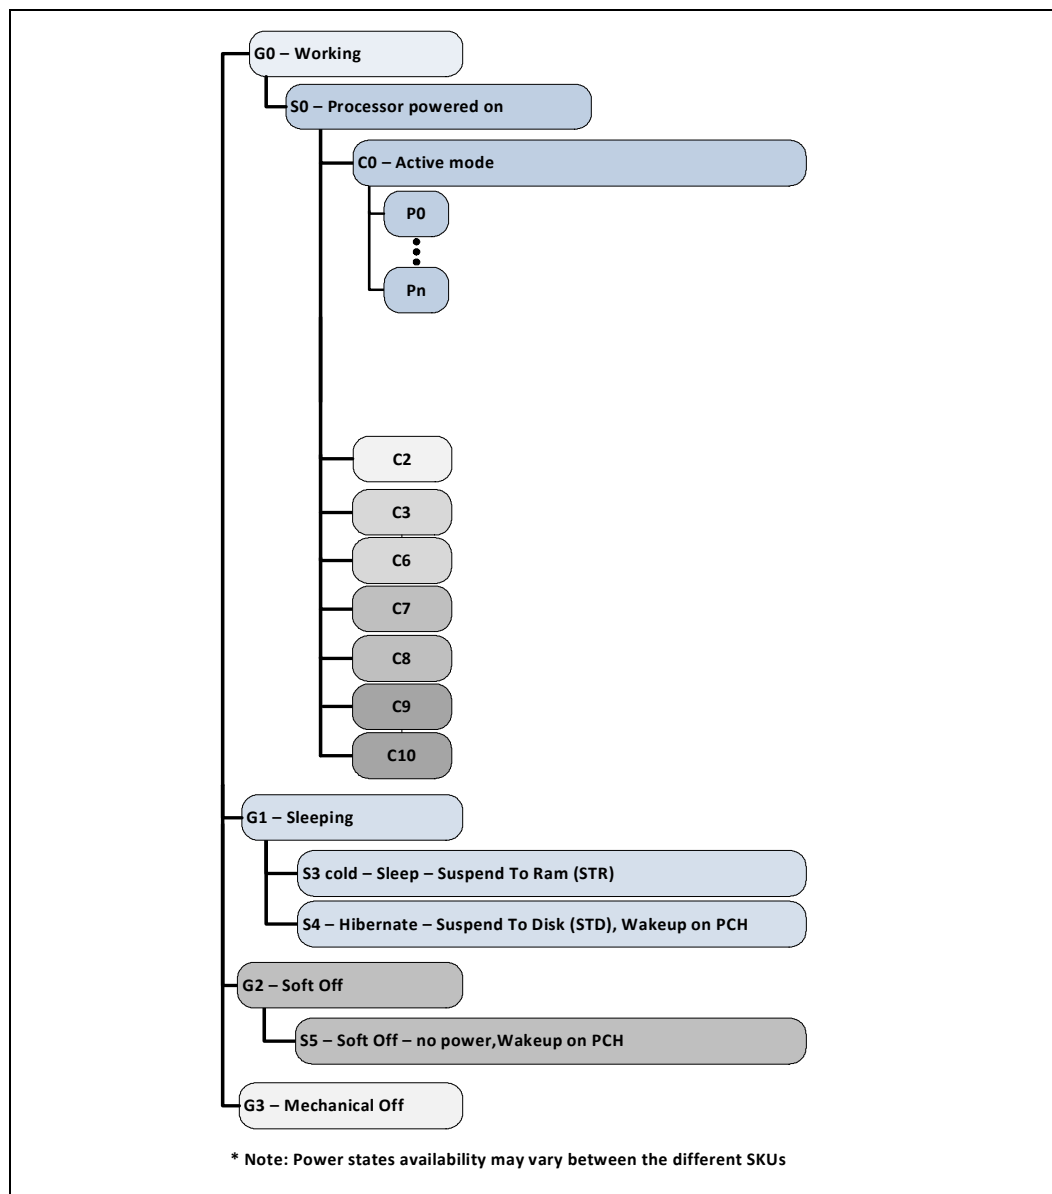


3 Power Management

This chapter provides information on the following power management topics:

- Advanced Configuration and Power Interface (ACPI) States
- Processor IA Core Power Management
- Integrated Memory Controller (IMC) Power Management
- Processor Graphics Power Management

Figure 3-1. Processor Power States





3.1 Advanced Configuration and Power Interface (ACPI) States Supported

This section describes the ACPI states supported by the processor.

Table 3-1. System States

State	Description
G0/S0/C0	Full On: CPU operating. Individual devices may be shut to save power. The different CPU operating levels are defined by Cx states.
G0/S0/Cx	Cx state: CPU manages C-states itself and can be in low power state
G1/S3	Suspend-To-RAM (STR): The system context is maintained in system DRAM, but power is shut to non-critical circuits. Memory is retained, and refreshes continue. All external clocks shut off; RTC clock and internal ring oscillator clocks are still toggling. In S3, SLP_S3 signal stays asserted, SLP_S4 and SLP_S5 are inactive until a wake occurs.
G1/S4	Suspend-To-Disk (STD): The context of the system is maintained on the disk. All power is then shut to the system except for the logic required to resume. Externally appears same as S5, but may have different wake events. In S4, SLP_S3 and SLP_S4 both stay asserted and SLP_S5 is inactive until a wake occurs.
G2/S5	Soft Off: System context not maintained. All power is shut except for the logic required to restart. A full boot is required when waking. Here, SLP_S3, SLP_S4 and SLP_S5 are all Active until a wake occurs.
G3	Mechanical OFF: System context not maintained. All power shut except for the RTC. No "Wake" events are possible, because the system does not have any power. This state occurs if the user removes the batteries, turns off a mechanical switch, or if the system power supply is at a level that is insufficient to power the "waking" logic. When system power returns, transition will depend on the state just prior to the entry to G3.

Table 3-2. Integrated Memory Controller (IMC) States

State	Description
Power up	CKE asserted. Active mode.
Pre-charge Power down	CKE de-asserted (not self-refresh) with all banks closed.
Active Power down	CKE de-asserted (not self-refresh) with minimum one bank active.
Self-Refresh	CKE de-asserted using device self-refresh.

Table 3-3. G, S, and C Interface State Combinations

Global (G) State	Sleep (S) State	Processor Package (C) State	Processor State	System Clocks	Description
G0	S0	C0	Full On	On	Full On
G0	S0	C2	Deep Sleep	On	Deep Sleep
G0	S0	C3	Deep Sleep	On	Deep Sleep
G0	S0	C6/C7	Deep Power Down	On	Deep Power Down
G0	S0	C8/C9/C10	Off	On	Deeper Power Down
G1	S3	Power off	Off	Off, except RTC	Suspend to RAM
G1	S4	Power off	Off	Off, except RTC	Suspend to Disk
G2	S5	Power off	Off	Off, except RTC	Soft Off
G3	N/A	Power off	Off	Power off	Hard off



3.2 Processor IA Core Power Management

While executing code, Enhanced Intel SpeedStep® Technology and Intel Speed Shift® technology optimizes the processor’s IA core frequency and voltage based on workload. Each frequency and voltage operating point is defined by ACPI as a P-state. When the processor is not executing code, it is idle. A low-power idle state is defined by ACPI as a C-state. In general, deeper power C-states have longer entry and exit latencies.

3.2.1 OS/HW Controlled P-states

3.2.1.1 Enhanced Intel SpeedStep® Technology

Enhanced Intel® SpeedStep® Technology enables OS to control and select P-state. For more information, refer to [Section 2.4.7, “Enhanced Intel SpeedStep® Technology”](#).

3.2.1.2 Intel® Speed Shift Technology

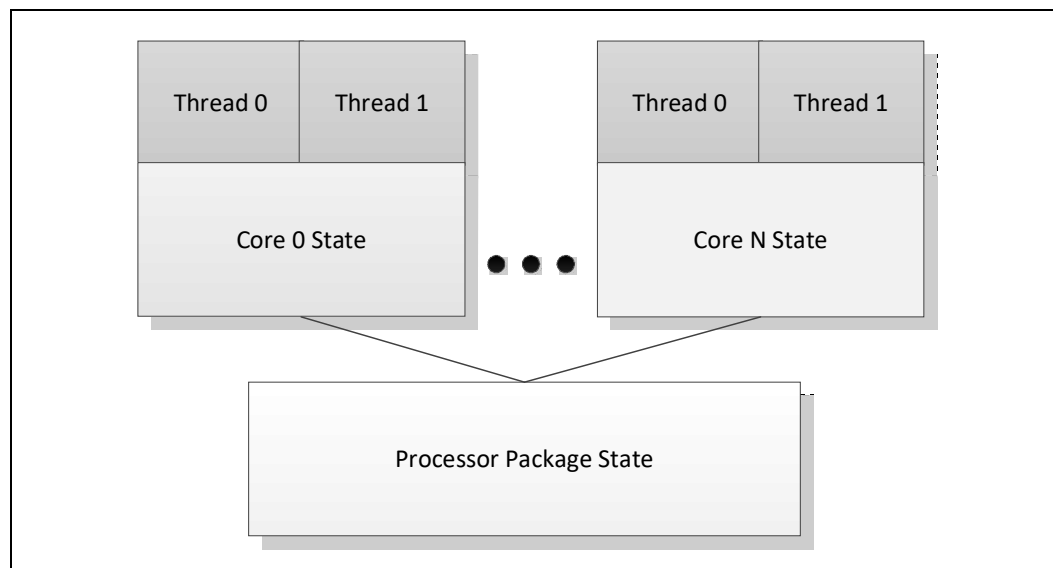
Intel® Speed Shift Technology is an energy efficient method of frequency control by the hardware rather than relying on OS control. For more details, refer to [Section 2.4.8, “Intel® Speed Shift Technology”](#).

3.2.2 Low-Power Idle States

When the processor is idle, low-power idle states (C-states) are used to save power. More power savings actions are taken for numerically higher C-states. However, deeper C-states have longer exit and entry latencies. Resolution of C-states occur at the thread, processor IA core, and processor package level. Thread-level C-states are available if Intel® Hyper-Threading Technology is enabled.

Caution: Long term reliability cannot be assured unless all the Low-Power Idle States are enabled.

Figure 3-2. Idle Power Management Breakdown of the Processor IA Cores





While individual threads can request low-power C-states, power saving actions only take place once the processor IA core C-state is resolved. processor IA core C-states are automatically resolved by the processor. For thread and processor IA core C-states, a transition to and from C0 state is required before entering any other C-state.

3.2.3 Requesting Low-Power Idle States

The primary software interfaces for requesting low-power idle states are through the MWAIT instruction with sub-state hints and the HLT instruction (for C1 and C1E). However, software may make C-state requests using the legacy method of I/O reads from the ACPI-defined processor clock control registers, referred to as P_LVLx. This method of requesting C-states provides legacy support for operating systems that initiate C-state transitions using I/O reads.

For legacy operating systems, P_LVLx I/O reads are converted within the processor to the equivalent MWAIT C-state request. Therefore, P_LVLx reads do not directly result in I/O reads to the system. The feature, known as I/O MWAIT redirection, should be enabled in the BIOS.

The BIOS can write to the C-state range field to restrict the range of I/O addresses that are trapped and emulate MWAIT like functionality. Any P_LVLx reads outside of this range do not cause an I/O redirection to MWAIT(Cx) like request. They fall through like a normal I/O instruction.

When P_LVLx I/O instructions are used, MWAIT sub-states cannot be defined. The MWAIT sub-state is always zero if I/O MWAIT redirection is used. By default, P_LVLx I/O redirections enable the MWAIT 'break on EFLAGS.IF' feature that triggers a wake up on an interrupt, even if interrupts are masked by EFLAGS.IF.

3.2.4 Processor IA Core C-State Rules

The following are general rules for all processor IA core C-states, unless specified otherwise:

- A processor IA core C-State is determined by the lowest numerical thread state (such as Thread 0 requests C1E while Thread 1 requests C6 state, resulting in a processor IA core C1E state). Refer the *G, S, and C Interface State Combinations* table.
- A processor IA core transitions to C0 state when:
 - An interrupt occurs
 - There is an access to the monitored address if the state was entered using an MWAIT/Timed MWAIT instruction
 - The deadline corresponding to the Timed MWAIT instruction expires
- An interrupt directed toward a single thread wakes up only that thread.
- If any thread in a processor IA core is active (in C0 state), the core's C-state will resolve to C0.
- Any interrupt coming into the processor package may wake any processor IA core.
- A system reset re-initializes all processor IA cores.



Table 3-4. Core C-states

Core C-State	C-State Request Instruction	Description
C0	N/A	The normal operating state of a processor IA core where code is being executed
C1	MWAIT(C1)	AutoHalt - core execution stopped, autonomous clock gating (package in C0 state)
C1E	MWAIT(C1E)	Core C1 + lowest frequency and voltage operating point (package in C0 state)
C6-C10	MWAIT(C6/7/7s/ C8/9/10) or IO read=P_LVL3/4/5/ 6/7/8	Processor IA, flush their L1 instruction cache, L1 data cache, and L2 cache to the LLC shared cache cores save their architectural state to a SRAM before reducing IA cores voltage, if possible may also be reduced to 0V. Core clocks are off.

Core C-State Auto-Demotion

In general, deeper C-states, such as C6 or C7, have long latencies and have higher energy entry/exit costs. The resulting performance and energy penalties become significant when the entry/exit frequency of a deeper C-state is high. Therefore, incorrect or inefficient usage of deeper C-states have a negative impact on battery life and idle power. To increase residency and improve battery life and idle power in deeper C-states, the processor supports C-state auto-demotion.

C-State auto-demotion:

- C7/C6 to C1/C1E

The decision to demote a processor IA core from C6/C7 to C1/C1E is based on each processor IA core's immediate residency history. Upon each processor IA core C6/C7 request, the processor IA core C-state is demoted to C1 until a sufficient amount of residency has been established. At that point, a processor IA core is allowed to go into C6 or C7. If the interrupt rate experienced on a processor IA core is high and the processor IA core is rarely in a deep C-state between such interrupts, the processor IA core can be demoted to a C1 state.

This feature is disabled by default. BIOS should enable it in the PMG_CST_CONFIG_CONTROL register. The auto-demotion policy is also configured by this register.

3.2.5 Package C-States

The processor supports C0, C2, C3, C6, C7, C8, C9, and C10 package states. The following is a summary of the general rules for package C-state entry. These apply to all package C-states, unless specified otherwise:

- A package C-state request is determined by the lowest numerical processor IA core C-state amongst all processor IA cores.
- A package C-state is automatically resolved by the processor depending on the processor IA core idle power states and the status of the platform components.
 - Each processor IA core can be at a lower idle power state than the package if the platform does not grant the processor permission to enter a requested package C-state.
 - The platform may allow additional power savings to be realized in the processor.

- For package C-states, the processor is not required to enter C0 before entering any other C-state.
- Entry into a package C-state may be subject to auto-demotion – that is, the processor may keep the package in a deeper package C-state than requested by the operating system if the processor determines, using heuristics, that the deeper C-state results in better power/performance.

The processor exits a package C-state when a break event is detected. Depending on the type of break event, the processor does the following:

- If a processor IA core break event is received, the target processor IA core is activated and the break event message is forwarded to the target processor IA core.
 - If the break event is not masked, the target processor IA core enters the processor IA core C0 state and the processor enters package C0.
 - If the break event is masked, the processor attempts to re-enter its previous package state.
- If the break event was due to a memory access or snoop request,
 - But the platform did not request to keep the processor in a higher package C-state, the package returns to its previous C-state.
 - And the platform requests a higher power C-state, the memory access or snoop request is serviced and the package remains in the higher power C-state.

Figure 3-3. Package C-State Entry and Exit

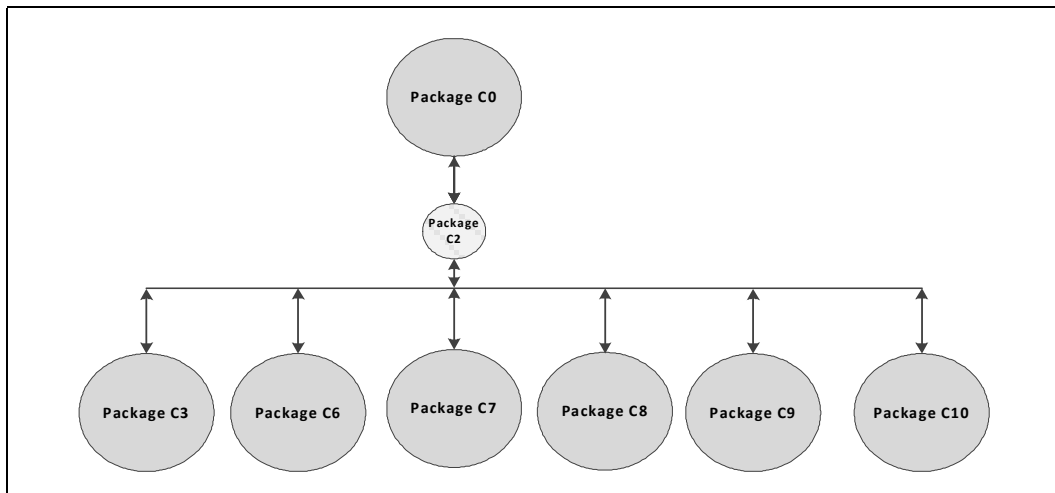


Table 3-5. Package C-States (Sheet 1 of 2)

Package C state	Description	Dependencies
C0	Processor active state	-



Table 3-5. Package C-States (Sheet 2 of 2)

Package C state	Description	Dependencies
C2	Cannot be requested explicitly by the Software. Memory path may be open. The processor will enter Package C2 when: <ul style="list-style-type: none"> • Transitioning from Package C0 to deep Package C state or from deep Package C state to Package C0. • All IA cores requested C6 or deeper + Processor Graphic cores in RC6 but there are constraints (LTR, programmed timer events in the near future and so forth) prevent entry to any state deeper than C2 state. • All IA cores requested C6 or deeper + Processor Graphic cores in RC6 but a device memory access request is received. Upon completion of all outstanding memory requests, the processor transitions back into a deeper package C-state. 	All processor IA cores in C6 or deeper. Processor Graphic cores in RC6.
C3	The processor will enter Package C3 when: <ul style="list-style-type: none"> • All IA cores in C6 or deeper + Processor Graphic cores in RC6. • The platform components/devices allows proper LTR for entering Package C3. 	All processor IA cores in C6 or deeper. Processor Graphics in RC6. memory in self refresh, memory clock stopped. LLC may be flushed and turned off
C6	The processor will enter Package C6 when: <ul style="list-style-type: none"> • All IA cores in C6 or deeper + Processor Graphic cores in RC6. • The platform components/devices allow proper LTR for entering Package C6. 	Package C3. BCLK is off. IMVP VRs voltage reduction/PSx state is possible.
C7	The processor will enter Package C7 when: <ul style="list-style-type: none"> • All IA cores in C7 or deeper + Processor Graphic cores in RC6. • The platform components/devices allow proper LTR for entering Package C7. 	Package C6. If all IA cores requested C7. LLC ways may be flushed until it is cleared. If the entire LLC is flushed, voltage will be removed from the LLC.
C7S	The processor will enter Package C7 when: <ul style="list-style-type: none"> • All IA cores in C7S or deeper + Processor Graphic cores in RC6. • The platform components/devices allow proper LTR for entering Package C7S. 	Package C6 If all IA cores requested C7S, LLC is flushed in a single step, voltage will be removed from the LLC.
C8	The processor will enter Package C8 when: <ul style="list-style-type: none"> • All IA cores in C8 or deeper + Processor Graphic cores in RC6. • The platform components/devices allow proper LTR for entering Package C8. 	Package C7 + LLC should be flushed at once.
C9	The processor will enter Package C9 when: <ul style="list-style-type: none"> • All IA cores in C9 or deeper + Processor Graphic cores in RC6. • The platform components/devices allow proper LTR for entering Package C9. 	Package C8. All IA cores in C9 or deeper. Display in PSR or powered off ¹ . VCCIO stays on.
C10	The processor will enter Package C10 when: <ul style="list-style-type: none"> • All IA cores in C10 + Processor Graphic cores in RC6. • The platform components/devices allow proper LTR for entering Package C10. 	Package C9. All VRs at PS4 or LPM. Crystal clock off. TCSS may enter lowest power state (TC cold) ²
Notes: 1. Display In PSR is only on single embedded panel configuration and panel support PSR feature. 2. At Package C10, TCSS can enter TC-cold when no device attached to any of TCSS ports.		



Package C-State Auto-Demotion

The Processor may demote the Package C state to a shallower C state, for example instead of going into package C10, it will demote to package C8 (and so on as required). The processor decision to demote the package C state is based on the required C states latencies, entry/exit energy/power and devices LTR.

Modern Standby

Modern Standby is a platform state. On display time out the OS requests the processor to enter package C10 and platform devices at RTD3 (or disabled) in order to attain low power in idle. Modern Standby requires proper BIOS and OS configuration.

Dynamic LLC Sizing

When all processor IA cores request C7 or deeper C-state, internal heuristics dynamically flushes the LLC. Once the processor IA cores enter a deep C-state, depending on their MWAIT sub-state request, the LLC is either gradually flushed N-ways at a time or flushed all at once. Upon the processor IA cores exiting to C0 state, the LLC is gradually expanded based on internal heuristics.

3.2.6 Package C-States and Display Resolutions

The integrated graphics engine has the frame buffer located in system memory. When the display is updated, the graphics engine fetches display data from system memory. Different screen resolutions and refresh rates have different memory latency requirements. These requirements may limit the deepest Package C-state the processor can enter. Other elements that may affect the deepest Package C-state available are the following:

- Display is on or off
- Single or multiple displays
- Native or non-native resolution
- Panel Self Refresh (PSR) technology

Note:

Display resolution is not the only factor influencing the deepest Package C-state the processor can get into. Device latencies, interrupt response latencies, and core C-states are among other factors that influence the final package C-state the processor can enter.

The following table lists display resolutions and deepest available package C-State. The display resolutions are examples using common values for blanking and pixel rate. Actual results will vary. The table shows the deepest possible Package C-state. System workload, system idle, and AC or DC power also affect the deepest possible Package C-state.

Table 3-6. Deepest Package C-State Available (Sheet 1 of 2)

		Y/U Processor Line ^{1,2}	
Resolution	Number of Displays	PSR Enabled	PSR Disabled
800x600 60 Hz	Single	PC10	PC8
1024x768 60 Hz	Single	PC10	PC8
1280x1024 60 Hz	Single	PC10	PC8


Table 3-6. Deepest Package C-State Available (Sheet 2 of 2)

		Y/U Processor Line ^{1,2}	
1920x1080 60 Hz	Single	PC10	PC8
1920x1200 60 Hz	Single	PC10	PC8
1920x1440 60 Hz	Single	PC10	PC8
2048x1536 60 Hz	Single	PC10	PC8
2560x1600 60 Hz	Single	PC10	PC8
2560x1920 60 Hz	Single	PC10	PC8
2880x1620 60 Hz	Single	PC10	PC8
2880x1800 60 Hz	Single	PC10	PC8
3200x1800 60 Hz	Single	PC10	PC8
3200x2000 60 Hz	Single	PC10	PC8
3840x2160 60 Hz	Single	PC10	PC8
4096x2160 60 Hz	Single	PC10	PC8
5120x3200 60 Hz	Single	PC10	PC8
Notes: 1. All Deep states are with Display ON. 2. The deepest C-state has variance, dependent various parameters such SW and Platform devices.			

3.3 Processor Graphics Power Management

3.3.1 Memory Power Savings Technologies

3.3.1.1 Intel® Rapid Memory Power Management (Intel® RMPM)

Intel® Rapid Memory Power Management (Intel® RMPM) conditionally places memory into self-refresh when the processor is in package C3 or deeper power state to allow the system to remain in the deeper power states longer for memory not reserved for graphics memory. Intel® RMPM functionality depends on graphics/display state (relevant only when processor graphics is being used), as well as memory traffic patterns generated by other connected I/O devices.

3.3.2 Display Power Savings Technologies

3.3.2.1 Intel® Seamless Display Refresh Rate Switching Technology (Intel® SDRRS Technology) with eDP* Port

Intel® DRRS provides a mechanism where the monitor is placed in a slower refresh rate (the rate at which the display is updated). The system is smart enough to know that the user is not displaying either 3D or media like a movie where specific refresh rates are required. The technology is very useful in an environment such as a plane where the user is in battery mode doing E-mail, or other standard office applications. It is also useful where the user may be viewing web pages or social media sites while in battery mode.



3.3.2.2 Intel® Automatic Display Brightness

Intel® Automatic Display Brightness feature dynamically adjusts the back-light brightness based upon the current ambient light environment. This feature requires an additional sensor to be on the panel front. The sensor receives the changing ambient light conditions and sends the interrupts to the Intel Graphics driver. As per the change in Lux, (current ambient light luminance), the new back-light setting can be adjusted through BLC (Back Light Control). The converse applies for a brightly lit environment. Intel® Automatic Display Brightness increases the back-light setting.

3.3.2.3 Smooth Brightness

The Smooth Brightness feature is the ability to make fine grained changes to the screen brightness. All Windows* 8 system that support brightness control are required to support Smooth Brightness control and it should be supporting 101 levels of brightness control. Apart from the Graphics driver changes, there may be few System BIOS changes required to make this feature functional.

3.3.2.4 Intel® Display Power Saving Technology (Intel® DPST) 6.3

The Intel® DPST technique achieves back-light power savings while maintaining a good visual experience. This is accomplished by adaptively enhancing the displayed image while decreasing the back-light brightness simultaneously. The goal of this technique is to provide equivalent end-user-perceived image quality at a decreased back-light power level.

1. The original (input) image produced by the operating system or application is analyzed by the Intel® DPST subsystem. An interrupt to Intel® DPST software is generated whenever a meaningful change in the image attributes is detected. (A meaningful change is when the Intel® DPST software algorithm determines that enough brightness, contrast, or color change has occurred to the displaying images that the image enhancement and back-light control needs to be altered.)
2. Intel® DPST subsystem applies an image-specific enhancement to increase image contrast, brightness, and other attributes.
3. A corresponding decrease to the back-light brightness is applied simultaneously to produce an image with similar user-perceived quality (such as brightness) as the original image.

Intel® DPST 6.3 has improved power savings without adversely affecting the performance.

3.3.2.5 Panel Self-Refresh 2 (PSR 2)

Panel Self-Refresh feature allows the Processor Graphics core to enter low-power state when the frame buffer content is not changing constantly. This feature is available on panels capable of supporting Panel Self-Refresh. Apart from being able to support, the eDP* panel should be eDP 1.4 compliant. PSR 2 adds partial frame updates and requires an eDP* 1.4 compliant panel.



3.3.2.6 Low-Power Single Pipe (LPSP)

Low-power single pipe is a power conservation feature that helps save power by keeping the inactive pipes powered OFF. This feature is enabled only in a single display configuration without any scaling functionalities. This feature is supported from 4th Generation Intel® Core™ processor family onwards. LPSP is achieved by keeping a single pipe enabled during eDP* only with minimal display pipeline support. This feature is panel independent and works with any eDP* panel (port A) in single display mode.

3.3.2.7 Intel® Smart 2D Display Technology (Intel® S2DDT)

Intel® S2DDT reduces display refresh memory traffic by reducing memory reads required for display refresh. Power consumption is reduced by less accesses to the IMC. Intel® S2DDT is only enabled in single pipe mode.

Intel® S2DDT is most effective with:

- Display images well suited to compression, such as text windows, slide shows, and so on. Poor examples are 3D games.
- Static screens such as screens with significant portions of the background showing 2D applications, processor benchmarks, and so on, or conditions when the processor is idle. Poor examples are full-screen 3D games and benchmarks that flip the display image at or near display refresh rates.

3.3.3 Processor Graphics Core Power Savings Technologies

3.3.3.1 Intel® Graphics Dynamic Frequency

Intel® Turbo Boost Technology 2.0 is the ability of the processor IA cores and graphics (Graphics Dynamic Frequency) cores to opportunistically increase frequency and/or voltage above the guaranteed processor and graphics frequency for the given part. Intel® Graphics Dynamic Frequency is a performance feature that makes use of unused package power and thermals to increase application performance. The increase in frequency is determined by how much power and thermal budget is available in the package, and the application demand for additional processor or graphics performance. The processor IA core control is maintained by an embedded controller. The graphics driver dynamically adjusts between P-States to maintain optimal performance, power, and thermals. The graphics driver will always place the graphics engine in its lowest possible P-State. Intel® Graphics Dynamic Frequency requires BIOS support.

3.3.3.2 Intel® Graphics Render Standby Technology (Intel® GRST)

Intel® Graphics Render Standby Technology is a technique designed to optimize the average power of the graphics part. The Graphics Render engine will be put in a sleep state, or Render Standby (RS), during times of inactivity or basic video modes. While in Render Standby state, the graphics part will place the VR (Voltage Regulator) into a low voltage state. Hardware will save the render context to the allocated context buffer when entering RS state and restore the render context upon exiting RS state.



3.3.3.3 Dynamic FPS (DFPS)

Dynamic FPS (DFPS) or dynamic frame-rate control is a runtime feature for improving power-efficiency for 3D workloads. Its purpose is to limit the frame-rate of full screen 3D applications without compromising on user experience. By limiting the frame rate, the load on the graphics engine is reduced, giving an opportunity to run the Processor Graphics at lower speeds, resulting in power savings. This feature works in both AC/DC modes.

3.4 System Agent Enhanced Intel SpeedStep® Technology

System Agent Enhanced Intel SpeedStep® Technology is a dynamic voltage frequency scaling of the System Agent clock based on memory utilization. Unlike processor core and package Enhanced Intel SpeedStep® Technology, System Agent Enhanced Intel SpeedStep® Technology has three valid operating points. When running light workload and SA Enhanced Intel SpeedStep® Technology is enabled, the DDR data rate may change as follows:

Before changing the DDR data rate, the processor sets DDR to self-refresh and changes the needed parameters. The DDR voltage remains stable and unchanged.

BIOS/MRC DDR training at max, mid and min frequencies sets I/O and timing parameters.

Refer to [Table 5-5, "SA Speed Enhanced Speed Steps \(SA-GV\) and Gear Mode Frequencies"](#).

3.5 Voltage Optimization

Voltage Optimization opportunistically provides reduction in power consumption, i.e., a boost in performance at a given PL1. Over time the benefit is reduced. There is no change to base frequency or turbo frequency. During system validation and tuning, this feature should be disabled to reflect processor power and performance that is expected over time.

3.6 ROP (Rest Of Platform) PMIC

In addition to discrete voltage regulators, Intel supports specific PMIC (Power Management Integrated Circuit) models to power the ROP rails. PMICs are typically classified as "Premium" or "Volume" ROP PMICs.

Note: Intel supports ROP PMIC as part of Y/U-Processor Lines.





4 Thermal Management

4.1 Y/U-Processor Line Thermal and Power Specifications

The following notes apply to Table 4-1, "TDP Specifications (U/Y-Processor Line)", Table 4-2, "Package Turbo Specifications", and Table 4-3, "Junction Temperature Specifications"

Note	Definition
1	The TDP and Configurable TDP values are the average power dissipation in junction temperature operating condition limit, for the SKU Segment and Configuration, for which the processor is validated during manufacturing when executing an associated Intel-specified high-complexity workload at the processor IA core frequency corresponding to the configuration and SKU.
2	TDP workload may consist of a combination of processor IA core intensive and graphics core intensive applications.
3	Can be modified at runtime, with MMIO and with PECI commands.
4	'Turbo Time Parameter' is a mathematical parameter (units of seconds) that controls the processor turbo algorithm using a moving average of energy usage. Do not set the Turbo Time Parameter to a value less than 0.1 seconds. refer to Section 4.2.1.2, "Platform Power Control" for further information.
5	Shown limit is a time averaged power, based upon the Turbo Time Parameter. Absolute product power may exceed the set limits for short durations or under virus or uncharacterized workloads.
6	Processor will be controlled to specified power limit as described in Section 2.4.6.1, "Intel® Turbo Boost Technology 2.0 Power Monitoring" . If the power value and/or 'Turbo Time Parameter' is changed during runtime, it may take a short period of time (approximately 3 to 5 times the 'Turbo Time Parameter') for the algorithm to settle at the new control limits.
7	This is a hardware default setting and not a behavioral characteristic of the part.
8	For controllable turbo workloads, the PL2 limit may be exceeded for up to 10 ms.
9	LPM power level is an opportunistic power and is not a guaranteed value as usages and implementations may vary.
10	Power limits may vary depending on if the product supports the 'TDP-up' and/or 'TDP-down' modes.
11	The processor die do not reach maximum sustained power simultaneously since the sum of the 2 die's estimated power budget is controlled to be equal to or less than the package TDP (PL1) limit
12	cTDP down power is based on GT2 equivalent graphics configuration. cTDP down does not decrease the number of active Processor Graphics EUs, but relies on Power Budget Management (PL1) to achieve the specified power level.
13	May vary based on SKU.
14	The formula of $PL2=PL1*1.25$ is the hardware default but may not represent the optimum value for processor performance. By including the benefits available from power and thermal management features the recommended value for PL2 can be higher
15	TDP workload does not reflect various I/O connectivity cases such as Thunderbolt.
16	Hardware default of PL1 Tau=1 s, By including the benefits available from power and thermal management features the recommended is to use PL1 Tau=28 s.



Table 4-1. TDP Specifications (U/Y-Processor Line)

Segment and Package	Processor IA Cores, Graphics Configuration and TDP	Configuration	Processor IA Core Frequency	Graphics core Frequency	Thermal Design Power (TDP) [w]	Notes	
U-Processor Line	4- Core 15W	Configurable TDP-Up	1.2 GHz to 1.5 GHz	1.05 GHz to 1.1 GHz	25	1	
		Base	1 GHz to 1.3 GHz		15		
		Configurable TDP-Down	0.7 GHz to 1.0 GHz		12/13^		
		LFM	400 MHz	300 MHz	N/A		
	2- Core 15W	Configurable TDP-Up	N/A	0.9 GHz	N/A		12/13^
		Base	1.2GHz		15		
		Configurable TDP-Down	0.9 GHz		N/A		
		LFM	400 MHz	300 MHz	N/A		
Y-Processor Line	4- Core 9W	Configurable TDP-Up	1.0 GHz to 1.3 GHz	1.05 GHz to 1.1 GHz	12	1	
		Base	0.7 GHz to 1.0 GHz		9		
		Configurable TDP-Down	N/A		N/A		
		LFM	400 MHz	300 MHz	N/A		
	2- Core 9W	Base	1.1 GHz	0.9 GHz	9	1	
		Configurable TDP-Down	0.8 GHz		8		
		LFM	400 MHz	300 MHz	N/A		
Note:							
1. The ~ sign stands for approximation and ^ stands for SKU dependent							

Table 4-2. Package Turbo Specifications

Segment and Package	Processor IA Cores, Graphics Configuration and TDP	Parameter	Minimum	Hardware Default	Maximum	Units	Notes
U-Processor Line	4/2- Core GT2 15W	Power Limit 1 Time (PL1 Tau)	0.01	1	448	S	1,2,3
		Power Limit 1 (PL1)	N/A	15	N/A	W	
		Power Limit 2 (PL2)	N/A	PL2=PL1*1.25	N/A	W	
Y-Processor Line	4- Core GT2 9W	Power Limit 1 Time (PL1 Tau)	0.01	1	448	S	1,2,3
		Power Limit 1 (PL1)	N/A	9	N/A	W	
		Power Limit 2 (PL2)	N/A	PL2=PL1*1.25	N/A	W	

**Table 4-3. Junction Temperature Specifications**

Segment	Symbol	Package Turbo Parameter	Temperature Range		TDP Specification Temperature Range		Units	Notes
			Minimum	Maximum	Minimum	Maximum		
U-Processor Line BGA	T_j	Junction temperature limit	0	100	35	100		1, 2
Y-Processor Line BGA	T_j	Junction temperature limit	0	100	0	90	°C	1, 2, 3

Notes:

- The thermal solution needs to ensure that the processor temperature does not exceed the TDP Specification Temperature.
- The processor junction temperature is monitored by Digital Temperature Sensors (DTS). For DTS accuracy, refer to [Section 4.2.3.2.1, "Digital Thermal Sensor Accuracy \(Taccuracy\)"](#).
- For the Y Processor Line to be specification compliance to the 90°C TDP specification temperature, TCC Offset = 10 and Tau value should be programmed. The recommended TCC_Offset averaging Tau value is 5 s. Ice Lake Processor External Design Specification Volume 2 of 2 (#571131)

4.2 Processor Thermal Management

The thermal solution provides both component-level and system-level thermal management. To allow optimal operation and long-term reliability of Intel processor-based systems, the system/processor thermal solution should be designed so that the processor:

- Remains below the maximum junction temperature (T_{jMAX}) specification at the maximum thermal design power (TDP).
- Conforms to system constraints, such as system acoustics, system skin-temperatures, and exhaust-temperature requirements.

Caution: Thermal specifications given in this chapter are on the component and package level and apply specifically to the processor. Operating the processor outside the specified limits may result in permanent damage to the processor and potentially other components in the system.

4.2.1 Thermal Considerations

The processor TDP is the maximum sustained power that should be used for design of the processor thermal solution. TDP is a power dissipation and junction temperature operating condition limit, specified in this document, that is validated during manufacturing for the base configuration when executing a near worst case commercially available workload as specified by Intel for the SKU segment. TDP may be exceeded for short periods of time or if running a very high power workload.

The processor integrates multiple processing IA cores, graphics cores and for some SKUs a PCH on a single package. This may result in power distribution differences across the package and should be considered when designing the thermal solution.

Intel® Turbo Boost Technology 2.0 allows processor IA cores to run faster than the base frequency. It is invoked opportunistically and automatically as long as the processor is conforming to its temperature, power delivery and current control limits. When Intel® Turbo Boost Technology 2.0 is enabled:

- Applications are expected to run closer to TDP more often as the processor will attempt to maximize performance by taking advantage of estimated available energy budget in the processor package.



- The processor may exceed the TDP for short durations to utilize any available thermal capacitance within the thermal solution. The duration and time of such operation can be limited by platform runtime configurable registers within the processor
- Graphics peak frequency operation is based on the assumption of only one of the graphics domains (GT/GTx) being active. This definition is similar to the IA core Turbo concept, where peak turbo frequency can be achieved when only one IA core is active. Depending on the workload being applied and the distribution across the graphics domains the user may not observe peak graphics frequency for a given workload or benchmark.
- Thermal solutions and platform cooling that are designed to less than thermal design guidance may experience thermal and performance issues.

4.2.1.1 Package Power Control

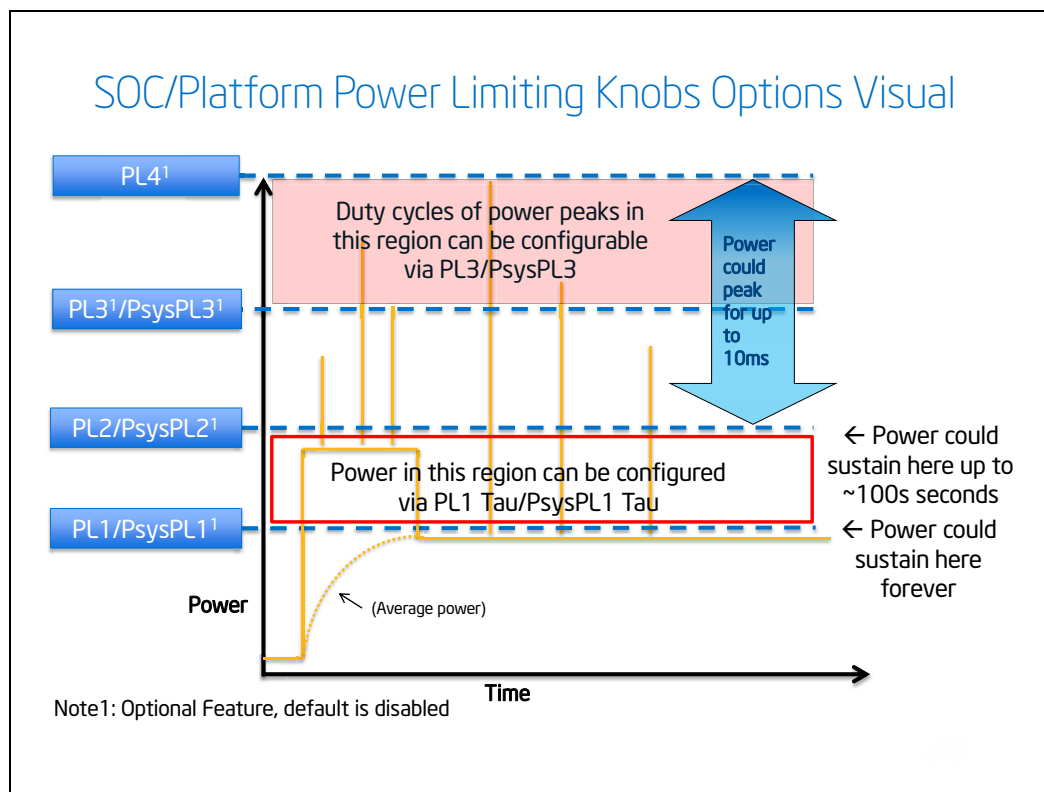
The package power control settings of PL1, PL2, PL3, PL4 and Tau allow the designer to configure Intel® Turbo Boost Technology 2.0 to match the platform power delivery and package thermal solution limitations.

- **Power Limit 1 (PL1):** A threshold for average power that will not exceed - recommend to set to equal TDP power. PL1 should not be set higher than thermal solution cooling limits.
- **Power Limit 2 (PL2):** A threshold that if exceeded, the PL2 rapid power limiting algorithms will attempt to limit the spike above PL2.
- **Power Limit 3 (PL3):** A threshold that if exceeded, the PL3 rapid power limiting algorithms will attempt to limit the duty cycle of spikes above PL3 by reactively limiting frequency. This is an optional setting
- **Power Limit 4 (PL4):** A limit that will not be exceeded, the PL4 power limiting algorithms will preemptively limit frequency to prevent spikes above PL4.
- **Turbo Time Parameter (Tau):** An averaging constant used for PL1 exponential weighted moving average (EWMA) power calculation.

Notes:

1. Implementation of Intel® Turbo Boost Technology 2.0 only requires configuring PL1, PL1 Tau and PL2.
2. PL3 and PL4 are disabled by default.

Figure 4-1. Package Power Control



4.2.1.2 Platform Power Control

The processor introduces Psys (Platform Power) to enhance processor power management. The Psys signal needs to be sourced from a compatible charger circuit and routed to the IMVP9 (voltage regulator). This signal will provide the total thermally relevant platform power consumption (processor and rest of platform) via SVID to the processor.

When the Psys signal is properly implemented, the system designer can utilize the package power control settings of PsysPL1/Tau, PsysPL2 and PsysPL3 for additional manageability to match the platform power delivery and platform thermal solution limitations for Intel® Turbo Boost Technology 2.0. The operation of the PsysPL1/tau, PsysPL2 and PsysPL3 is analogous to the processor power limits described in [Section 4.2.1.1, "Intel® Turbo Boost Technology 2.0 availability may vary between the different SKUs. Package Power Control"](#).

- **Platform Power Limit 1 (PsysPL1):** A threshold for average platform power that will not be exceeded - recommend to set to equal platform thermal capability.
- **Platform Power Limit 2 (PsysPL2):** A threshold that if exceeded, the PsysPL2 rapid power limiting algorithms will attempt to limit the spikes above PsysPL2.
- **Platform Power Limit 3 (PsysPL3):** A threshold that if exceeded, the PsysPL3 rapid power limiting algorithms will attempt to limit the duty cycle of spikes above PsysPL3 by reactively limiting frequency.
- **PsysPL1 Tau:** An averaging constant used for PsysPL1 exponential weighted moving average (EWMA) power calculation.



- The Psys signal and associated power limits / Tau are optional for the system designer and disabled by default.
- The Psys data will not include power consumption for charging.

4.2.1.3 Turbo Time Parameter (Tau)

Turbo Time Parameter (Tau) is a mathematical parameter (units of seconds) that controls the Intel® Turbo Boost Technology 2.0 algorithm. During a maximum power turbo event, the processor could sustain PL2 for a duration longer than the Turbo Time Parameter. If the power value and/or Turbo Time Parameter is changed during runtime, it may take some time based on the new Turbo Time Parameter level for the algorithm to settle at the new control limits. The time varies depending on the magnitude of the change, power limits and other factors. There is an individual Turbo Time Parameter associated with Package Power Control and Platform Power Control.

4.2.2 Configurable TDP (cTDP) and Low-Power Mode

Configurable TDP (cTDP) and Low-Power Mode (LPM) form a design option where the processor's behavior and package TDP are dynamically adjusted to a desired system performance and power envelope. Configurable TDP and Low-Power Mode technologies offer opportunities to differentiate system design while running active workloads on select processor SKUs through scalability, configuration and adaptability. The scenarios or methods by which each technology is used are customizable but typically involve changes to PL1 and associated frequencies for the scenario with a resultant change in performance depending on system's usage. Either technology can be triggered by (but are not limited to) changes in OS power policies or hardware events such as docking a system, flipping a switch or pressing a button. cTDP and LPM are designed to be configured dynamically and do not require an operating system reboot.

Note: Configurable TDP and Low-Power Mode technologies are not battery life improvement technologies.

4.2.2.1 Configurable TDP

Note: Configurable TDP availability may vary between the different SKUs.

With cTDP, the processor is now capable of altering the maximum sustained power with an alternate processor IA core base frequency. Configurable TDP allows operation in situations where extra cooling is available or situations where a cooler and quieter mode of operation is desired.

cTDP consists of three modes as shown in the following table.

Table 4-4. Configurable TDP Modes (Sheet 1 of 2)

Mode	Description
Base	The average power dissipation and junction temperature operating condition limit, specified in Table 4-1, "TDP Specifications (U/Y-Processor Line)" and Table 4-3, "Junction Temperature Specifications" for the SKU Segment and Configuration, for which the processor is validated during manufacturing when executing an associated Intel-specified high-complexity workload at the processor IA core frequency corresponding to the configuration and SKU.



Table 4-4. Configurable TDP Modes (Sheet 2 of 2)

Mode	Description
TDP-Up	The SKU-specific processor IA core frequency where manufacturing confirms logical functionality within the set of operating condition limits specified for the SKU segment and Configurable TDP-Up configuration in Table 4-1, "TDP Specifications (U/Y-Processor Line)" and Table 4-3, "Junction Temperature Specifications" . The Configurable TDP-Up Frequency and corresponding TDP is higher than the processor IA core Base Frequency and SKU Segment Base TDP.
TDP-Down	The processor IA core frequency where manufacturing confirms logical functionality within the set of operating condition limits specified for the SKU segment and Configurable TDP-Down configuration in Table 4-1, "TDP Specifications (U/Y-Processor Line)" and Table 4-3, "Junction Temperature Specifications" . The Configurable TDP-Down Frequency and corresponding TDP is lower than the processor IA core Base Frequency and SKU Segment Base TDP.

In each mode, the Intel® Turbo Boost Technology 2.0 power limits are reprogrammed along with a new OS controlled frequency range. The cTDP mode does not change the maximum per-processor IA core turbo frequency.

4.2.2.2 Low-Power Mode

Low-Power Mode (LPM) can provide cooler and quieter system operation. By combining several active power limiting techniques, the processor can consume less power while running at equivalent low frequencies. Active power is defined as processor power consumed while a workload is running and does not refer to the power consumed during idle modes of operation.

LPM can be configured to use each of the following methods to reduce active power:

- Restricting package power control limits and Intel® Turbo Boost Technology availability
- Off-Lining processor IA core activity (Move processor traffic to a subset of cores)
- Placing a processor IA Core at LFM or LSF (Lowest Supported Frequency)
- Utilizing IA clock modulation
- Reducing number of active EUs to GT2 equivalent (applicable for GT3 SKUs Only)
- LPM power as listed in the *TDP Specifications* table is defined at point which processor IA core working at LSF, GT = RPn and 1 IA core active

Off-lining processor IA core activity is the ability to dynamically scale a workload to a limited subset of cores in conjunction with a lower turbo power limit. It is one of the main vectors available to reduce active power. However, not all processor activity is ensured to be able to shift to a subset of cores. Shifting a workload to a limited subset of cores allows other processor IA cores to remain idle and save power. Therefore, when LPM is enabled, less power is consumed at equivalent frequencies.

Minimum Frequency Mode (MFM) of operation, which is the Lowest Supported Frequency (LSF) at the LFM voltage, has been made available for use under LPM for further reduction in active power beyond LFM capability to enable cooler and quieter modes of operation.

4.2.3 Thermal Management Features

Occasionally the processor may operate in conditions that are near to its maximum operating temperature. This can be due to internal overheating or overheating within the platform. In order to protect the processor and the platform from thermal failure, several thermal management features exist to reduce package power consumption and thereby temperature in order to remain within normal operating limits. Furthermore, the processor supports several methods to reduce memory power.

4.2.3.1 Adaptive Thermal Monitor

The purpose of the Adaptive Thermal Monitor is to reduce processor IA core power consumption and temperature until it operates below its maximum operating temperature. Processor IA core power reduction is achieved by:

- Adjusting the operating frequency (using the processor IA core ratio multiplier) and voltage.
- Modulating (starting and stopping) the internal processor IA core clocks (duty cycle).

The Adaptive Thermal Monitor can be activated when the package temperature, monitored by any Digital Thermal Sensor (DTS), meets its maximum operating temperature. The maximum operating temperature implies maximum junction temperature T_{jMAX} .

Reaching the maximum operating temperature activates the Thermal Control Circuit (TCC). When activated the TCC causes both the processor IA core and graphics core to reduce frequency and voltage adaptively. The Adaptive Thermal Monitor will remain active as long as the package temperature remains at its specified limit. Therefore, the Adaptive Thermal Monitor will continue to reduce the package frequency and voltage until the TCC is de-activated.

T_{jMAX} is factory calibrated and is not user configurable. The default value is software visible.

The Adaptive Thermal Monitor does not require any additional hardware, software drivers, or interrupt handling routines. It is not intended as a mechanism to maintain processor thermal control to PL1 = TDP. The system design should provide a thermal solution that can maintain normal operation when PL1 = TDP within the intended usage range.

Adaptive Thermal Monitor protection is always enabled.

4.2.3.1.1 TCC Activation Offset

TCC Activation Offset can be set as an offset from T_{jMAX} to lower the onset of TCC and Adaptive Thermal Monitor. In addition, there is an optional time window (τ) to manage processor performance at the TCC Activation offset value via an EWMA (Exponential Weighted Moving Average) of temperature.

TCC Activation Offset with $\tau=0$

An offset (degrees Celsius) can be written, the offset value will be subtracted from the value found in bits [23:16]. When the time window (τ) is set to zero, there will be no averaging, the offset, will be subtracted from the T_{jMAX} value and used as a new



max temperature set point for Adaptive Thermal Monitoring. This will have the same behavior as in prior products to have TCC activation and Adaptive Thermal Monitor to occur at this lower target silicon temperature.

If enabled, the offset should be set lower than any other passive protection such as ACPI _PSV trip points

TCC Activation Offset with Tau

To manage the processor with the EWMA (Exponential Weighted Moving Average) of temperature, an offset (degrees Celsius) is written, and the time window (Tau) is written. The Offset value will be subtracted from the value found in bits [23:16] and be the temperature.

The processor will manage to this average temperature by adjusting the frequency of the various domains. The instantaneous T_j can briefly exceed the average temperature. The magnitude and duration of the overshoot is managed by the time window value (Tau).

This averaged temperature thermal management mechanism is in addition, and not instead of $T_{j_{MAX}}$ thermal management. That is, whether the TCC activation offset is 0 or not, TCC Activation will occur at $T_{j_{MAX}}$.

4.2.3.1.2 Frequency / Voltage Control

Upon Adaptive Thermal Monitor activation, the processor attempts to dynamically reduce processor temperature by lowering the frequency and voltage operating point. The operating points are automatically calculated by the processor IA core itself and do not require the BIOS to program them as with previous generations of Intel processors. The processor IA core will scale the operating points such that:

- The voltage will be optimized according to the temperature, the processor IA core bus ratio and number of processor IA cores in deep C-states.
- The processor IA core power and temperature are reduced while minimizing performance degradation.

Once the temperature has dropped below the trigger temperature, the operating frequency and voltage will transition back to the normal system operating point.

Once a target frequency/bus ratio is resolved, the processor IA core will transition to the new target automatically.

- On an upward operating point transition the voltage transition precedes the frequency transition.
- On a downward transition the frequency transition precedes the voltage transition.
- The processor continues to execute instructions. However, the processor will halt instruction execution for frequency transitions.

If a processor load-based Enhanced Intel SpeedStep® Technology/P-state transition is initiated while the Adaptive Thermal Monitor is active, there are two possible outcomes:

- If the P-state target frequency is higher than the processor IA core optimized target frequency, the P-state transition will be deferred until the thermal event has been completed.
- If the P-state target frequency is lower than the processor IA core optimized target frequency, the processor will transition to the P-state operating point.



4.2.3.1.3 Clock Modulation

If the frequency/voltage changes are unable to end an Adaptive Thermal Monitor event, the Adaptive Thermal Monitor will utilize clock modulation. Clock modulation is done by alternately turning the clocks off and on at a duty cycle (ratio between clock "on" time and total time) specific to the processor. The duty cycle is factory configured to 25% on and 75% off and cannot be modified. The period of the duty cycle is configured to 32 microseconds when the Adaptive Thermal Monitor is active. Cycle times are independent of processor frequency. A small amount of hysteresis has been included to prevent excessive clock modulation when the processor temperature is near its maximum operating temperature. Once the temperature has dropped below the maximum operating temperature, and the hysteresis timer has expired, the Adaptive Thermal Monitor goes inactive and clock modulation ceases. Clock modulation is automatically engaged as part of the Adaptive Thermal Monitor activation when the frequency/voltage targets are at their minimum settings. Processor performance will be decreased when clock modulation is active. Snooping and interrupt processing are performed in the normal manner while the Adaptive Thermal Monitor is active.

Clock modulation will not be activated by the Package average temperature control mechanism.

4.2.3.1.4 TT2/TT1 (Thermal Throttling Point)

As the processor approaches TJMax a throttling mechanisms will engage to protect the processor from over-heating and provide control thermal budgets. Achieving this is done by reducing IA and other subsystem agent's voltages and frequencies in a gradual and coordinated manner that varies depending on the dynamics of the situation. IA frequencies and voltages will be directed down as low as LFM (Lowest Frequency Mode). Further restricts are possible via Thermal Trolling point (TT1) under conditions where thermal budget cannot be re-gained fast enough with voltages and frequencies reduction alone. TT1 keeps the same processor voltage and clock frequencies the same yet skips clock edges to produce effectively slower clocking rates. This will effectively result in observed frequencies below LFM on the Windows PERF monitor.

4.2.3.2 Digital Thermal Sensor

Each processor has multiple on-die Digital Thermal Sensor (DTS) that detects the processor IA, GT and other areas of interest instantaneous temperature.

Temperature values from the DTS can be retrieved through:

- A software interface using processor internal identifiers.
- A processor hardware interface.

When temperature is retrieved by the processor internal identifiers, it is the instantaneous temperature of the given DTS. When temperature is retrieved using PECCI, it is the average of the highest DTS temperature in the package over a 256 ms time window. Intel recommends using the PECCI reported temperature for platform thermal control that benefits from averaging, such as fan speed control. The average DTS temperature may not be a good indicator of package Adaptive Thermal Monitor activation or rapid increases in temperature that triggers the Out of Specification status bit.

Code execution is halted in C1 or deeper C- states. Package temperature can still be monitored through PECCI in lower C-states*.



Unlike traditional thermal devices, the DTS outputs a temperature relative to the maximum supported operating temperature of the processor ($T_{j_{MAX}}$), regardless of TCC activation offset. It is the responsibility of software to convert the relative temperature to an absolute temperature. The absolute reference temperature is readable. The temperature returned by the DTS is an implied negative integer indicating the relative offset from $T_{j_{MAX}}$. The DTS does not report temperatures greater than $T_{j_{MAX}}$. The DTS-relative temperature readout directly impacts the Adaptive Thermal Monitor trigger point. When a package DTS indicates that it has reached the TCC activation (a reading of 0x0, except when the TCC activation offset is changed), the TCC will activate and indicate an Adaptive Thermal Monitor event. A TCC activation will lower both processor IA core and graphics core frequency, voltage, or both. Changes to the temperature can be detected using two programmable thresholds located in the processor thermal internal identifiers. These thresholds have the capability of generating interrupts using the processor IA core's local APIC.

4.2.3.2.1 Digital Thermal Sensor Accuracy (Taccuracy)

The error associated with DTS measurements will not exceed ± 5 °C within the entire operating range.

4.2.3.2.2 Fan Speed Control with Digital Thermal Sensor

Digital Thermal Sensor based fan speed control (T_{FAN}) is a recommended feature to achieve optimal thermal performance. At the T_{FAN} temperature, Intel recommends full cooling capability before the DTS reading reaches $T_{j_{MAX}}$.

4.2.3.3 PROCHOT# Signal

PROCHOT# (processor hot) is asserted by the processor when the TCC is active. Only a single PROCHOT# pin exists at a package level. When any DTS temperature reaches the TCC activation temperature, the PROCHOT# signal will be asserted. PROCHOT# assertion policies are independent of Adaptive Thermal Monitor enabling. PROCHOT# signal can be configured as:

Input only: PROCHOT is driven by an external device.

Output only: PROCHOT is driven by processor.

Bi-Directional: Both Processor and external device can drive PROCHOT signal.

4.2.3.4 PROCHOT Input Only

It is recommended to set by default the PROCHOT# signal to **input only**. Processor is monitoring only the PROCHOT# assertions and not the PROCHOT# level. PROCHOT# maximum toggling frequency should not exceed 10Khz.

When PROCHOT is set to Input only two features are enabled:

- **Fast PROCHOT:** Activate up to 10 uS after PROCHOT assertion and reduce the processor frequency by half.
- **PROCHOT Demotion Algorithm:** designed to improve system performance during multiple PROCHOT assertions (detailed explanation in [Section 4.2.3.7, "PROCHOT Demotion Algorithm"](#))



4.2.3.5 PROCHOT Output Only

Legacy state, PROCHOT is driven by the processor to external device.

4.2.3.6 Bi-Directional PROCHOT#

When configured as an input or bi-directional signal, PROCHOT# can be used for thermally protecting other platform components should they overheat as well. Processor is monitoring only the PROCHOT# assertions and not the PROCHOT# level. PROCHOT# maximum toggling frequency should not exceed 10 kHz

When PROCHOT# is driven by an external device:

- The package will immediately transition to the lowest P-State (Pn) supported by the processor IA cores and graphics cores (LFM). This is contrary to the internally-generated Adaptive Thermal Monitor response.
- Clock modulation is not activated.

The processor package will remain at the lowest supported P-state until the system de-asserts PROCHOT#. The processor can be configured to generate an interrupt upon assertion and de-assertion of the PROCHOT# signal.


When PROCHOT# is configured as a bi-directional signal and PROCHOT# is asserted by the processor, it is impossible for the processor to detect a system assertion of PROCHOT#. The system assertion will have to wait until the processor de-asserts PROCHOT# before PROCHOT# action can occur due to the system assertion. While the processor is hot and asserting PROCHOT#, the power is reduced but the reduction rate is slower than the system PROCHOT# response of < 100 us. The processor thermal control is staged in smaller increments over many milliseconds. This may cause several milliseconds of delay to a system assertion of PROCHOT# while the output function is asserted.

4.2.3.7 PROCHOT Demotion Algorithm

PROCHOT demotion algorithm designed to improve system performance following multiple EC PROCHOT consecutive assertions. During each PROCHOT assertion processor will immediately transition to the lowest P-State (Pn) supported by the processor IA cores and graphics cores (LFM). When detecting several PROCHOT consecutive assertions the processor will reduce the max frequency in order to reduce the PROCHOT assertions events. The processor will keep reducing the frequency until no consecutive assertions detected. The processor will raise the frequency if no consecutive PROCHOT assertion events will occur. PROCHOT demotion algorithm enabled only when the PROCHOT is configured as input.



Figure 4-2. PROCHOT Demotion Signal Description



4.2.3.8 Voltage Regulator Protection using PROCHOT#

PROCHOT# may be used for thermal protection of voltage regulators (VR). System designers can create a circuit to monitor the VR temperature and assert PROCHOT# and, if enabled, activate the TCC when the temperature limit of the VR is reached. When PROCHOT# is configured as a bi-directional or input only signal, if the system assertion of PROCHOT# is recognized by the processor, it will result in an immediate transition to the lowest P-State (Pn) supported by the processor IA cores and graphics cores. Systems should still provide proper cooling for the VR and rely on bi-directional PROCHOT# only as a backup in case of system cooling failure. Overall, the system thermal design should allow the power delivery circuitry to operate within its temperature specification even while the processor is operating at its TDP.


4.2.3.9 Thermal Solution Design and PROCHOT# Behavior

With a properly designed and characterized thermal solution, it is anticipated that PROCHOT# will only be asserted for very short periods of time when running the most power intensive applications. The processor performance impact due to these brief periods of TCC activation is expected to be so minor that it would be immeasurable. However, an under-designed thermal solution that is not able to prevent excessive assertion of PROCHOT# in the anticipated ambient environment may:

- Cause a noticeable performance loss.
- Result in prolonged operation at or above the specified maximum junction temperature and affect the long-term reliability of the processor.
- May be incapable of cooling the processor even when the TCC is active continuously (in extreme situations).

4.2.3.10 Low-Power States and PROCHOT# Behavior

Depending on package power levels during package C-states, outbound PROCHOT# may de-assert while the processor is idle as power is removed from the signal. Upon wake up, if the processor is still hot, the PROCHOT# will re-assert, although typically package idle state residency should resolve any thermal issues. The PECCI interface is fully operational during all C-states and it is expected that the platform continues to manage processor IA core and package thermals even during idle states by regularly polling for thermal data over PECCI.



4.2.3.11 THRMTRIP# Signal

Regardless of enabling the automatic or on-demand modes, in the event of a catastrophic cooling failure, the package will automatically shut down when the silicon has reached an elevated temperature that risks physical damage to the product. At this point the THRMTRIP# signal will go active.

4.2.3.12 Critical Temperature Detection

Critical Temperature detection is performed by monitoring the package temperature. This feature is intended for graceful shutdown before the THRMTRIP# is activated. However, the processor execution is not guaranteed between critical temperature and THRMTRIP#. If the Adaptive Thermal Monitor is triggered and the temperature remains high, a critical temperature status and sticky bit are latched and the condition also generates a thermal interrupt, if enabled.

4.2.3.13 On-Demand Mode

The processor provides an auxiliary mechanism that allows system software to force the processor to reduce its power consumption using clock modulation. This mechanism is referred to as "On-Demand" mode and is distinct from Adaptive Thermal Monitor and bi-directional PROCHOT#. The processor platforms should not rely on software usage of this mechanism to limit the processor temperature. On-Demand Mode can be accomplished using processor internal identifiers or chipset I/O emulation. On-Demand Mode may be used in conjunction with the Adaptive Thermal Monitor. However, if the system software tries to enable On-Demand mode at the same time the TCC is engaged, the factory configured duty cycle of the TCC will override the duty cycle selected by the On-Demand mode.

4.2.3.14 I/O Emulation-Based On-Demand Mode

I/O emulation-based clock modulation provides legacy support for operating system software that initiates clock modulation through I/O writes to ACPI defined processor clock control registers on the chipset (PROC_CNT). Thermal throttling using this method will modulate all processor IA cores simultaneously.

4.2.4 Intel® Memory Thermal Management

The processor provides thermal protection for system memory by throttling memory traffic when using either DIMM modules or a memory down implementation. Two levels of throttling are supported by the processor, either a warm threshold or hot threshold that is customizable through memory mapped I/O registers. Throttling based on the warm threshold should be an intermediate level of throttling. Throttling based on the hot threshold should be the most severe. The amount of throttling is dynamically controlled by the processor.

The on Die Thermal Sensor (ODTS) uses a physical thermal sensor on DRAM dies. ODTS is available for DDR4 and LPDDR4/x. It is used to set refresh rate according to DRAM temperature. The memory controller reads LPDDR4/x MR4 or DDR4 MR3 and configures the DDR refresh rate accordingly.





5 Memory

5.1 System Memory Interface

5.1.1 Processor SKU Support Matrix

Table 5-1. DDR Support Matrix Table

Technology	DDR4	LPDDR4/x
Processor	U	U/Y
Maximum Frequency [MT/s]	3200	3733
VDDQ ⁵ [V]	1.2	1.1
VDD2 ⁶ [V]	1.2	1.1
Channels x Bits	2 x 64	4 x 32
DPC ¹	1	-
RPC ²	2	2
Die Density [Gb]	8,16	4,8
Ballmap Mode	IL ³ /NIL	NIL
Notes: <ol style="list-style-type: none"> 1DPC refer to when only 1DIMM slot per channel is routed. RPC = Rank Per Channel. An Interleave SoDIMM/MD placements like butterfly or back-to-back supported with Non-Interleave ballmap mode at U Processor Line Memory down of all technologies should be implemented homogeneous means that all DRAM devices should be from the same vendor and have the same part number. Implementing a mix of DRAM devices may cause serious signal integrity and functional issues. There is no support for memory modules with different technologies or capacities on opposite sides of the same memory module. If one side of a memory module is populated, the other side is either identical or empty. LPDDR4/x Processor VDDQ is 1.1 V. LPDDR4 DRAM VDDQ voltage is 1.1 V, VDD2 is 1.1 V LPDDR4x DRAM VDDQ voltage is 0.6 V, VDD2 is 1.1 V 		

Table 5-2. DDR technology Support Matrix (Sheet 1 of 2)

Form Factor	Ball count	DDR4	LPDDR4	LPDDR4x
SODIMM	260	U	-	-
SODIMM + ECC	260	U	-	-
x16 SDP (1R) ¹	96	U	-	-
x16 DDP (1R) ^{1,2}	96	U	-	-
x8 SDP (1R) ¹	78	U	-	-
x32 (1R, 2R) ¹	200	-	Y, U	Y, U
x64 (1R, 2R) ^{1,3}	432	-	Y, U	Y, U
x64 (1R, 2R) ^{1,4}	556	-	-	Y



Table 5-2. DDR technology Support Matrix (Sheet 2 of 2)

Form Factor	Ball count	DDR4	LPDDR4	LPDDR4x
Note: 1. Memory down of all technologies should be implemented homogeneously, which means that all DRAM devices should be from the same vendor and have the same part number. Implementing a mix of DRAM devices may cause serious signal integrity and functional issues. 2. DDPx16 is pending on sample availability. 3. U Processor LP4/4x x64 topology is Non-POR topology. 4. Y Processor LP4/4x 556 ball topology is Non-POR topology.				

Table 5-3. DDR Max Capacity per System

Processor Line	LP4/x x32 (2x 8Gb)	DDR4 1DPC 8Gb	DDR4 1DPC 16Gb
U	32 GB	32 GB	64 GB
Y	32 GB	N/A	N/A

Table 5-4. LPDDR4/x Sub-Channels Population Rules

Number of DRAMs	DRAM Type	Sub-Channel Population
1	x32	N/A
2 ²	x32	DRAM 0 is connected to Sub Channel A ¹ DRAM 1 is connected to Sub Channel C ¹
3	x32	N/A
4	x32	DRAM 0 is connected to Sub Channel A DRAM 1 is connected to Sub Channel B DRAM 2 is connected to Sub Channel C DRAM 3 is connected to Sub Channel D
1	x64	DRAM 0 is connected to Sub Channel A and C ¹
2	x64	DRAM 0 is connected to Sub Channel A and C DRAM 1 is connected to Sub Channel B and D or DRAM 0 is connected to Sub Channel A and B DRAM 1 is connected to Sub Channel C and D
Note: 1. Connecting DRAM 0 to sub channel A and B accordingly is possible but less preferred as A and C are performance/bandwidth optimized. 2. To work with one x64 channel need to set "MemConfig->ForceSingleSubchannel = 1" in BIOS		

Table 5-5. SA Speed Enhanced Speed Steps (SA-GV) and Gear Mode Frequencies (Sheet 1 of 2)

Technology	DDR max rate [MT/s]	SAGV-Low DDR CLK, Gear	SAGV-High ³ DDR CLK, Gear	SAGV-Max BW DDR CLK, Gear
DDR4	2666	2133, G2	U - 2400, G1	2666, G2
	2933	2133, G2	U - 2400, G1	2933, G2
	3200	2133, G2	U - 2400, G1	3200, G2
LPDDR4/x	3200	2133, G2	Y - 3200, G2	3200, G2
			U - 2400, G1	
	3733	2133, G2	Y - 3200, G2	3733, G2
			U - 2400, G1	


Table 5-5. SA Speed Enhanced Speed Steps (SA-GV) and Gear Mode Frequencies (Sheet 2 of 2)

Technology	DDR max rate [MT/s]	SAGV-Low DDR CLK, Gear	SAGV-High ³ DDR CLK, Gear	SAGV-Max BW DDR CLK, Gear
Notes: 1. Processor supports dynamic gearing technology where the Memory Controller can run at 1:1 (Gear-1, Legacy mode) or 1:2 (Gear-2 mode) ratio of DRAM speed. Gear ratio is the ratio of DRAM speed to Memory Controller Clock. MC Channel Width equal to DDR Channel width multiply by Gear Ratio. 2. SA-GV modes a. Low - Low frequency point, Min Power point. Characterized by low power, low BW, high latency. System will stay at this point during low to moderate BW consumption. b. Mid - Max Bandwidths Point, this point is the max possible BW point, the DRAM freq limited by Silicon Configuration/BIOS/SPD. Characterized by moderate power and latency, high BW. This point intended for high GT and moderate-high IA BW c. High - High Point, the minimum memory latency point, Characterized by high power, low latency, moderate BW. Only during IA performance workloads the system will to switch to this point and only in case this point can provide enough BW. 3. High Point per SKU is optional support target for QS.				

Table 5-6. Supported DDR4 Non-ECC SODIMM Module Configurations (U-Processor Line)

Raw Card Version	Speed (MT/s)	DIMM Capacity	DRAM Device Technology	DRAM Organization	# of DRAM Devices	# of Ranks	# of Row/ Col Address Bits	# of Banks Inside DRAM	Page Size
A	3200	8 GB	8 GB	1024M x 8	8	1	16/10	16	8K
A	3200	16 GB	16 GB	2048M x 8	8	1	17/10	16	8K
C	3200	4 GB	8 GB	512M x 16	4	1	16/10	8	8K
C	3200	8 GB	16 GB	1024M x 16	4	1	17/10	8	8K
E	3200	16 GB	8 GB	1024M x 8	16	2	16/10	16	8K
E	3200	32 GB	16 GB	2048M x 8	16	2	17/10	16	8K

Table 5-7. Supported DDR4 Memory Down Device Configurations (U-Processor Line)

Max System Capacity ³	Speed (MT/s)	PKG Type (Die bits x Package bits)	DRAM Organization / Package Type	Package Density	Die Density	Dies Per Channel	Rank Per Channel	PKGs Per channel	Physical Device Rank	Banks Inside DRAM	Page Size
32GB	3200	SDP 8x8	1024Mx8	8 GB	8 GB	16	2	16	1	16	8K
64GB	3200	SDP 8x8	2048Mx8	16 GB	16 GB	16	2	16	1	16	8K
8GB	3200	SDP 16x16	512Mx16	8 GB	8 GB	4	1	4	1	8	8K
16GB ¹	3200	SDP 16x16	1024Mx16	16 GB	16 GB	4	1	4	1	8	8K
16GB	3200	DDP 8x16	1024Mx16	16 GB	8 GB	8	1	4	1	16	8K
32GB ^{2,3}	3200	DDP 8x16	2048Mx16	32 GB	16 GB	8	1	4	1	16	8K

Notes:

- For SDP: 1Rx16 using 16 GB die density - the maximum system capacity is 16 GB.
- For DDP: 1Rx16 using 16 GB die density - the maximum system capacity is 32 GB.
- Pending on sample availability.
- Max system capacity refer to system with two channels populated.



5.1.1.1 LPDDR4/x Supported Memory Modules and Devices

Table 5-8. Supported LPDDR4/x x32 DRAMs Configurations (Y/U-Processor Line)

Max System Capacity ⁴	PKG Type (Die bits per Ch x PKG bits) ²	Die Density per Channel	PKG Density	Rank Per PKGs
4 GB	DDP 16x32	4 GB	8 GB	1
8 GB	QDP 16x32	4 GB	16 GB	2
8 GB	DDP 16x32	8 GB	16 GB	1
16 GB	QDP 16x32	8 GB	32 GB	2
32GB	ODP 16x32 (Byte mode)	8 GB	64 GB	2

Notes:

- x32 BGA devices are 200 balls.
- DDP = Dual Die Package, QDP = Quad Die Package, ODP-Octal Die Package.
- Each LPDDR4 channel include two sub-channels.
- Max system capacity refer to system with all four sub-channels populated.

Table 5-9. Supported LPDDR4/x x64 DRAMs Configurations (U/Y-Processor Line)

Max System Capacity ⁴	PKG Type (Die bits per Ch x PKG bits) ²	Die Density per Channel	Ball Count per PKG	PKG Density	DRAM Channels per PKGs	Processor Line	Rank Per PKGs
8 GB	QDP 16x64	8 GB	432	32 GB	4	U/Y	1
16 GB	ODP 16x64	8 GB	432	64 GB	4	U/Y	2
8 GB ¹	QDP 16x64	8 GB	556	32 GB	4	Y	1
16 GB ¹	ODP 16x64	8 GB	556	64 GB	4	Y	2

Notes:

- Y Processor LP4/4x 556 ball topology is Non-POR topology.
- QDP = Quad Die Package, ODP-Octal Die Package.
- Each LPDDR4 channel include two sub-channels.
- Max system capacity refer to system with all four sub-channels populated.

5.1.2 System Memory Timing Support

The IMC supports the following DDR Speed Bin, CAS Write Latency (CWL), and command signal mode timings on the main memory interface:

- tCL = CAS Latency
- tRCD = Activate Command to READ or WRITE Command delay
- tRP = PRECHARGE Command Period
- tRPb = per-bank PRECHARGE time
- tRPab = all-bank PRECHARGE time
- CWL = CAS Write Latency
- Command Signal modes:
 - 1N indicates a new DDR4/LPDDR4 command may be issued every clock
 - 2N indicates a new DDR4 command may be issued every two clocks


Table 5-10. DDR4 System Memory Timing Support

DRAM Device	Transfer Rate (MT/s)	tCL (tCK)	tRCD (ns)	tRP (ns)	CWL (tCK)	DPC	CMD Mode
DDR4	3200	22	13.75	13.75	9,10,11,12,14,16,18,20	1	2N

Table 5-11. LPDDR4/x System Memory Timing Support

DRAM Device	mode	Transfer Rate (MT/s)	tCL (tCK)	tRCD (ns)	tRPpb (ns)	tRPab (ns)	WL (tCK) Set B
LPDDR4/x	x8	3733	36	18	18	21	30
	x16	3733	32	18	18	21	30

5.1.3 System Memory Controller Organization Modes

The IMC supports two memory organization modes, single-channel and dual-channel. Depending upon how the DDR Schema and DIMM Modules are populated in each memory channel, a number of different configurations can exist.

Single-Channel Mode

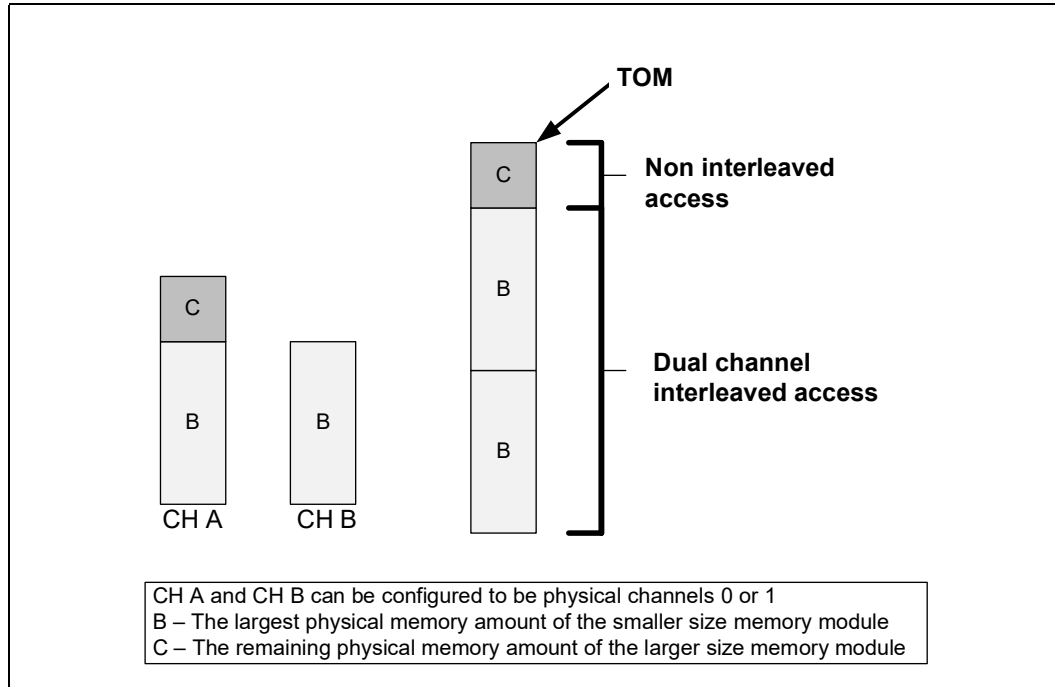
In this mode, all memory cycles are directed to a single channel. Single-Channel mode is used when either the Channel A or Channel B DIMM connectors are populated in any order, but not both.

Dual-Channel Mode – Intel® Flex Memory Technology Mode

The IMC supports Intel Flex Memory Technology Mode. Memory is divided into a symmetric and asymmetric zone. The symmetric zone starts at the lowest address in each channel and is contiguous until the asymmetric zone begins or until the top address of the channel with the smaller capacity is reached. In this mode, the system runs with one zone of dual-channel mode and one zone of single-channel mode, simultaneously, across the whole memory array.

Note: Channels A and B can be mapped for physical channel 0 and 1 respectively or vice versa; however, channel A size should be greater or equal to channel B size.

Figure 5-1. Intel® Flex Memory Technology Operations



Dual-Channel Symmetric Mode (Interleaved Mode)

Dual-Channel Symmetric mode, also known as interleaved mode, provides maximum performance on real world applications. Addresses are ping-ponged between the channels after each cache line (64-byte boundary). If there are two requests, and the second request is to an address on the opposite channel from the first, that request can be sent before data from the first request has returned. If two consecutive cache lines are requested, both may be retrieved simultaneously, since they are ensured to be on opposite channels. Use Dual-Channel Symmetric mode when both Channel A and Channel B DIMM connectors are populated in any order, with the total amount of memory in each channel being the same.

When both channels are populated with the same memory capacity and the boundary between the dual channel zone and the single channel zone is the top of memory, IMC operates completely in Dual-Channel Symmetric mode.

Note: The DRAM device technology and width may vary from one channel to the other.

5.1.4 System Memory Frequency

In all modes, the frequency of system memory is the lowest frequency of all memory modules placed in the system, as determined through the SPD registers on the memory modules. The system memory controller supports a single DIMM connector per channel. If DIMMs with different latency are populated across the channels, the BIOS will use the slower of the two latencies for both channels. For Dual-Channel modes both channels should have a DIMM connector populated. For Single-Channel mode, only a single channel can have a DIMM connector populated.



5.1.5 Technology Enhancements of Intel® Fast Memory Access (Intel® FMA)

The following sections describe the Just-in-Time Scheduling, Command Overlap, and Out-of-Order Scheduling Intel FMA technology enhancements.

Just-in-Time Command Scheduling

The memory controller has an advanced command scheduler where all pending requests are examined simultaneously to determine the most efficient request to be issued next. The most efficient request is picked from all pending requests and issued to system memory Just-in-Time to make optimal use of Command Overlapping. Thus, instead of having all memory access requests go individually through an arbitration mechanism forcing requests to be executed one at a time, they can be started without interfering with the current request allowing for concurrent issuing of requests. This allows for optimized bandwidth and reduced latency while maintaining appropriate command spacing to meet system memory protocol.

Command Overlap

Command Overlap allows the insertion of the DRAM commands between the Activate, Pre-charge, and Read/Write commands normally used, as long as the inserted commands do not affect the currently executing command. Multiple commands can be issued in an overlapping manner, increasing the efficiency of system memory protocol.

Out-of-Order Scheduling

While leveraging the Just-in-Time Scheduling and Command Overlap enhancements, the IMC continuously monitors pending requests to system memory for the best use of bandwidth and reduction of latency. If there are multiple requests to the same open page, these requests would be launched in a back to back manner to make optimum use of the open memory page. This ability to reorder requests on the fly allows the IMC to further reduce latency and increase bandwidth efficiency.

5.1.6 Data Scrambling

The system memory controller incorporates a Data Scrambling feature to minimize the impact of excessive di/dt on the platform system memory VRs due to successive 1s and 0s on the data bus. Past experience has demonstrated that traffic on the data bus is not random and can have energy concentrated at specific spectral harmonics creating high di/dt which is generally limited by data patterns that excite resonance between the package inductance and on die capacitances. As a result the system memory controller uses a data scrambling feature to create pseudo-random patterns on the system memory data bus to reduce the impact of any excessive di/dt.

5.1.7 Data Swapping

By default, the processor supports on-board data swapping in two manners (for all segments and DRAM technologies):

- Byte (8 DQ + DQS_N + DQS_P) swapping is allowed within a channel. For LPDDR4/x, Byte swapping is allowed within each 32-bit sub channel.
- Bit swapping is allowed within each Byte.



5.1.8 DDR I/O Interleaving

The processor supports I/O interleaving, which has the ability to swap DDR bytes for routing considerations. BIOS configures the I/O interleaving mode before DDR training.

Note: The Y/U-Processor line package is optimized only for Non-Interleaving mode (NIL).

There are two supported modes:

- Interleave (IL)
- Non-Interleave (NIL)


The following table and figure describe the pin mapping between the IL and NIL modes.

Table 5-12. Interleave (IL) and Non-Interleave (NIL) Modes Pin Mapping

IL (DDR4)		NIL (DDR4)		NIL (LPDDR4)	
Channel	Byte	Channel	Byte	Sub Channel	Byte
DDR0	Byte0	DDR0	Byte0	DDR_A	Byte0
DDR0	Byte1	DDR0	Byte2	DDR_A	Byte2
DDR0	Byte2	DDR0	Byte4	DDR_B	Byte0
DDR0	Byte3	DDR0	Byte6	DDR_B	Byte2
DDR0	Byte4	DDR1	Byte0	DDR_C	Byte0
DDR0	Byte5	DDR1	Byte2	DDR_C	Byte2
DDR0	Byte6	DDR1	Byte4	DDR_D	Byte0
DDR0	Byte7	DDR1	Byte6	DDR_D	Byte2
DDR1	Byte0	DDR0	Byte1	DDR_A	Byte1
DDR1	Byte1	DDR0	Byte3	DDR_A	Byte3
DDR1	Byte2	DDR0	Byte5	DDR_B	Byte1
DDR1	Byte3	DDR0	Byte7	DDR_B	Byte3
DDR1	Byte4	DDR1	Byte1	DDR_C	Byte1
DDR1	Byte5	DDR1	Byte3	DDR_C	Byte3
DDR1	Byte6	DDR1	Byte5	DDR_D	Byte1
DDR1	Byte7	DDR1	Byte7	DDR_D	Byte3

Notes: Y/U - supports NIL only.

Figure 5-2. Interleave (IL) and Non-Interleave (NIL) Modes Mapping



5.1.9 Data Swapping

By default, the processor supports on-board data swapping in two manners (for all segments and DRAM technologies):

- Byte (DQ+DQS) swapping between bytes in the same channel.
- Bit swapping within specific byte. ECC bits swap is allowed.

5.1.10 DRAM Clock Generation

Every supported rank has a differential clock pair. There are a total of four clock pairs driven directly by the processor to DRAM.

5.1.11 DRAM Reference Voltage Generation

The memory controller has the capability of generating the LPDDR4 and DDR4 Reference Voltage (VREF) internally for both read and write operations. The generated VREF can be changed in small steps, and an optimum VREF value is determined for both during a cold boot through advanced training procedures in order to provide the best voltage to achieve the best signal margins.

5.1.12 Data Swizzling

All Processor Lines does not have die-to-package DDR swizzling.

5.2 Integrated Memory Controller (IMC) Power Management

The main memory is power managed during normal operation and in low-power ACPI C-states.

5.2.1 Disabling Unused System Memory Outputs

Any system memory (SM) interface signal that goes to a memory in which it is not connected to any actual memory devices (such as SODIMM connector is unpopulated, or is single-sided) is tri-stated. The benefits of disabling unused SM signals are:

- Reduced power consumption.
- Reduced possible overshoot/undershoot signal quality issues seen by the processor I/O buffer receivers caused by reflections from potentially un-terminated transmission lines.
- When a given rank is not populated, the corresponding control signals (CLK_P/CLK_N/CKE/ODT/CS) are not driven.

At reset, all rows should be assumed to be populated, until it can be proven that they are not populated. This is due to the fact that when CKE is tri-stated with a DRAMs present, the DRAMs are not ensured to maintain data integrity. CKE tri-state should be enabled by BIOS where appropriate, since at reset all rows should be assumed to be populated.

5.2.2 DRAM Power Management and Initialization

The processor implements extensive support for power management on the memory interface. Each channel drives 4 CKE pins, one per rank.

The CKE is one of the power-saving means. When CKE is off, the internal DDR clock is disabled and the DDR power is reduced. The power-saving differs according to the selected mode and the DDR type used. For more information, refer to the IDD table in the DDR specification.

The processor supports four different types of power-down modes in package C0 state. The different power-down modes can be enabled through configuring PM PDWN config register. The type of CKE power-down can be configured through PDWN_mode (bits 15:12) and the idle timer can be configured through PDWN_idle_counter (bits 11:0).

The different power-down modes supported are:

- **No power-down** (CKE disable)
- **Active power-down (APD):** This mode is entered if there are open pages when de-asserting CKE. In this mode the open pages are retained. Power-saving in this mode is the lowest. Power consumption of DDR is defined by IDD3P. Exiting this mode is fined by tXP – small number of cycles. For this mode, DRAM DLL should be on.
- **PPD/DLL-off:** In this mode the data-in DLLs on DDR are off. Power-saving in this mode is the best among all power modes. Power consumption is defined by IDD2P. Exiting this mode is defined by tXP, but also tXPDLL (10–20 according to DDR type) cycles until first data transfer is allowed. For this mode, DRAM DLL should be off.
- **Precharged power-down (PPD):** This mode is entered if all banks in DDR are pre-charged when de-asserting CKE. Power-saving in this mode is intermediate – better than APD, but less than DLL-off. Power consumption is defined by IDD2P. Exiting this mode is defined by tXP. The difference from APD mode is that when waking-up, all page-buffers are empty.) The LPDDR does not have a DLL. As a result, the power savings are as good as PPD/DLL-off but will have lower exit latency and higher performance.



The CKE is determined per rank, whenever it is inactive. Each rank has an idle counter. The idle-counter starts counting as soon as the rank has no accesses, and if it expires, the rank may enter power-down while no new transactions to the rank arrives to queues. The idle-counter begins counting at the last incoming transaction arrival. It is important to understand that since the power-down decision is per rank, the IMC can find many opportunities to power down ranks, even while running memory intensive applications; the savings are significant (may be few Watts, according to DDR specification). This is significant when each channel is populated with more ranks.

Selection of power modes should be according to power-performance or thermal trade off of a given system:

- When trying to achieve maximum performance and power or thermal consideration is not an issue: use no power-down
- In a system which tries to minimize power-consumption, try using the deepest power-down mode possible – PPD/DLL-off with a low idle timer value
- In high-performance systems with dense packaging (that is, tricky thermal design) the power-down mode should be considered in order to reduce the heating and avoid DDR throttling caused by the heating.

The default value that BIOS configures in PM PDWN config register is 6080 – that is, PPD/DLL-off mode with idle timer of 0x80 (128 DCLKs). This is a balanced setting with deep power-down mode and moderate idle timer value.

The idle timer expiration count defines the # of DCLKs that a rank is idle that causes entry to the selected power mode. As this timer is set to a shorter time the IMC will have more opportunities to put the DDR in power-down. There is no BIOS hook to set this register. Customers choosing to change the value of this register can do it by changing it in the BIOS. For experiments, this register can be modified in real time if BIOS does not lock the IMC registers.

5.2.2.1 Initialization Role of CKE

During power-up, CKE is the only input to the SDRAM that has its level recognized (other than the reset pin) once power is applied. It should be driven LOW by the DDR controller to make sure the SDRAM components float DQ and DQS during power-up. CKE signals remain LOW (while any reset is active) until the BIOS writes to a configuration register. Using this method, CKE is ensured to remain inactive for much longer than the specified 200 micro-seconds after power and clocks to SDRAM devices are stable.

5.2.2.2 Conditional Self-Refresh

During S0 idle state, system memory may be conditionally placed into self-refresh state when the processor is in package C3 or deeper power state. Refer to [Section 3.3.1.1, “Intel® Rapid Memory Power Management \(Intel® RMPM\)”](#) for more details on conditional self-refresh with Intel HD Graphics enabled.

When entering the S3 – Suspend-to-RAM (STR) state or S0 conditional self-refresh, the processor IA core flushes pending cycles and then enters SDRAM ranks that are not used by the processor graphics into self-refresh. The CKE signals remain LOW so the SDRAM devices perform self-refresh.

The target behavior is to enter self-refresh for package C3 or deeper power states as long as there are no memory requests to service.

5.2.2.3 Dynamic Power-Down

Dynamic power-down of memory is employed during normal operation. Based on idle conditions, a given memory rank may be powered down. The IMC implements aggressive CKE control to dynamically put the DRAM devices in a power-down state. The processor IA core controller can be configured to put the devices in active powerdown (CKE de-assertion with open pages) or precharge power-down (CKE de-assertion with all pages closed). Precharge power-down provides greater power savings but has a bigger performance impact, since all pages will first be closed before putting the devices in power-down mode.

If dynamic power-down is enabled, all ranks are powered up before doing a refresh cycle and all ranks are powered down at the end of refresh.

5.2.2.4 DRAM I/O Power Management

Unused signals should be disabled to save power and reduce electromagnetic interference. This includes all signals associated with an unused memory channel. Clocks, CKE, ODT and CS signals are controlled per DIMM rank and will be powered down for unused ranks.

The I/O buffer for an unused signal should be tri-stated (output driver disabled), the input receiver (differential sense-amp) should be disabled, and any DLL circuitry related ONLY to unused signals should be disabled. The input path should be gated to prevent spurious results due to noise on the unused signals (typically handled automatically when input receiver is disabled).

5.2.3 DDR Electrical Power Gating

The DDR I/O of the processor supports Electrical Power Gating (DDR-EPG) while the processor is at C3 or deeper power state.

In C3 or deeper power state, the processor internally gates VDDQ for the majority of the logic to reduce idle power while keeping all critical DDR pins such as CKE and VREF in the appropriate state.

In C7 or deeper power state, the processor internally gates Vccio for all non-critical state to reduce idle power.

In S3 or C-state transitions, the DDR does not go through training mode and will restore the previous training information.

5.2.4 Power Training

BIOS MRC performing Power Training steps to reduce DDR I/O power while keeping reasonable operational margins still guaranteeing platform operation. The algorithms attempt to weaken ODT, driver strength and the related buffers parameters both on the MC and the DRAM side and find the best possible trade-off between the total I/O power and the operational margins using advanced mathematical models.





6 USB-C* Sub System

The USB protocol has five speeds: Low-speed, Full-speed, High-speed, SuperSpeed and Super-Speed plus. Refer to [Table 6-1, "USB Specifications"](#).

The USB-C* cables carry two physical buses, one for USB2 ("Low/Full/High" speeds) and one for the USB3 additions ("SuperSpeed/SuperSpeed+") the buses may be referred as "USB2" and "USB3" throughout this chapter.

Note: USB ports of processor implement USB3 and connect to the USB3 part of the USB-C* connector.

6.1 General Characteristics

- U processor supports a maximum of four USB-C* ports.
- Y processor supports a maximum of three USB-C* ports.
- xHCI (host controller) and xDCI (device controller) implemented in processor in addition to the controllers in the PCH and not replacement.
- No support for USB Type-A on the processor side, if needed those should be hanging of the PCH.

6.2 USB3.x Supported Features

- Support power saving when USB-C* disconnected.
- Host
 - USB3.x, SSIC (HSIC- USB2 is supplied via PCH xHCI).
 - Aggregate BW through the controller at least 3 GB/s, direct connection or over Thunderbolt.
 - At least one port of SSIC.
 - Wake capable on each host port from S0i3, Sx: Wake on Connects, Disconnects, Device Wake.
- Device
 - Aggregate BW ~ 1.2 GB/s.
 - D0i2 and D0i3 power gating.
 - Wake capable on host initiated wakes when system is in S0i3, Sx
 - Available on all ports.
- Port Routing Control for Dual Role Capability
 - Needs to support SW/FW and ID pin based control to detect host versus device attach.
 - SW mode requires PD controller or other FW to control.
- USB-R device to host controller connection is over UTMI+ links.



6.3 TCSS USB Blocks

The processor added xHCI/xDCI controllers (refer to [Section 6.3.1, "USB Controllers"](#)) for TCSS USB support. The native USB path proceed from the memory directly to PHY (refer to [Section 6.3.2, "PHY"](#)). In Thunderbolt™ mode, USB is encapsulated via Thunderbolt™ switch and sent via Thunderbolt™ protocol through PHY to USB-C* connector. Block diagram shows processor internal block diagram.

6.3.1 USB Controllers

Extensible Host Controller Interface (xHCI) is a the interface specification that defines Host Controller for Universal Serial bus (USB), which is capable of interfacing with USB 1.x, 2.0, and 3.x compatible devices. In case that a device (example, USB mouse) was connected to the computer, the computer will work as Host and the xHCI will be activated inside the CPU.

Extensible Device Controller Interface (xDCI) is a the interface specification that defines Device Controller for Universal Serial bus (USB), which is capable of interfacing with USB 1.x, 2.0, and 3.x compatible devices. In case that the computer is connected as a device (example, tablet connected to desktop) to other computer then the xDCI controller will be activated inside the device and will talk to the Host at the other computer.

Note: The processor USB subsystem incorporates a USB 3.0 device controller allows data transfers of up to 5 Gb/s and USB3.1 host controller that allows data transfers of up to 10 Gb/s. This controllers are instantiated in the processor die as a separate PCI function functionality for the USB-C* capable ports.

Table 6-1. USB Specifications

Protocol Name	Data Rate	USB3.0	USB3.1
Low - speed	1.5 Mbps	+	+
Full - speed	12 Mbps	+	+
High - speed	480 Mbps	+	+
SuperSpeed	5 Gbps	+	+
SuperSpeed+	10 Gbps	-	+
Note: USB2 ("Low/Full/High" speeds) implemented in PCH			

6.3.2 PHY

PHY is capable of supporting set of pins to be configured either as USB-C* connector pins or legacy DDI (DisplayPort/HDMI) connector pins.

Table 6-2. USB-C* Supported Configuration (Sheet 1 of 2)

Lane1	Lane2	Comments
Thunderbolt™	Thunderbolt™	Both lanes at the same speed, one of (20.6g/10.3g/20g/10g)
Thunderbolt™	No connect	20.6g/10.3g/20g/10g
No connect	Thunderbolt™	
USB3.1 Gen2	No connect	Any combination of USB3.1 and 3.0
No connect	USB 3.1 Gen2	
USB 3.1	DPx2	Any of HBR3/HBR2/HBR1/RBR for DP and USB3.1 Gen2
DPx2	USB3.1	



Table 6-2. USB-C* Supported Configuration (Sheet 2 of 2)

Lane1	Lane2	Comments
DPx4		Both lanes at the same DP rate - no support for 2x DPx2 USB-C connector

Table 6-3. USB-C* Non-Supported Configuration

Lane1	Lane2	Comments
#	PCIe* Gen3/2/1	No PCIe* native support
PCIe* Gen3/2/1	#	
#	Thunderbolt™	No support for Thunderbolt™ with any other protocol
Thunderbolt™	#	
USB 3.1	USB 3.1	No support for Multi-lane USB

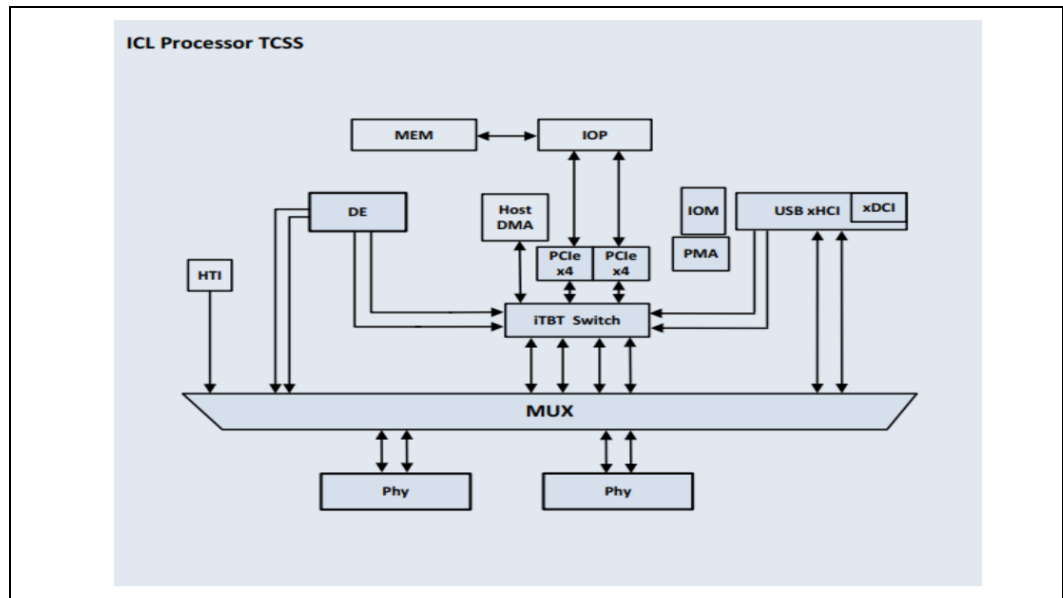
Table 6-4. PCIe* via TBT Configuration

TBT IPs	TBT_PCIE	U USB-C* Ports	Y USB-C* Ports
TBT_DMA0	TBT_PCIE0	TC0	TC0
	TBT_PCIE1	TC1	TC1
TBT_DMA1	TBT_PCIE2	TC2	TC2
	TBT_PCIE3	TC3	N/A

6.3.3 Integrated Thunderbolt™

For the Integrated Thunderbolt™ controller refer to [Chapter 7, “Thunderbolt™”](#).

Figure 6-1. USB-C* Sub-system Block Diagram





6.4 Power states

TCSS supports up to 4 Ports, each port supports low power state independently.

TCSS power management allow Processor power saving at following scenarios:

- Device attached without Traffic (Idle)
- Device attached, TCSS Controllers at D3
- No Device attached, TCSS Controllers at D3

Table 6-5. TCSS power states

TCSS Power State	Description	Allowed Package C States
TC0 Active	TCSS On, at least one TCSS controller is active.	PC0-PC3
TC0 Idle	All TCSS controllers are at Idle and accessible.	PC0-PC8
TC7	Deepest power state with Device attached. All TCSS controllers in D3	PC0-PC10
TC10	Deep power state with No Device attached. All TCSS controllers in D3	PC0-PC10
TC-Cold	Deepest power state with No Device attached. All TCSS controllers in D3 All TCSS Power rails turnoff, TCCold is triggered by the BIOS/OS	PC10





7 Thunderbolt™

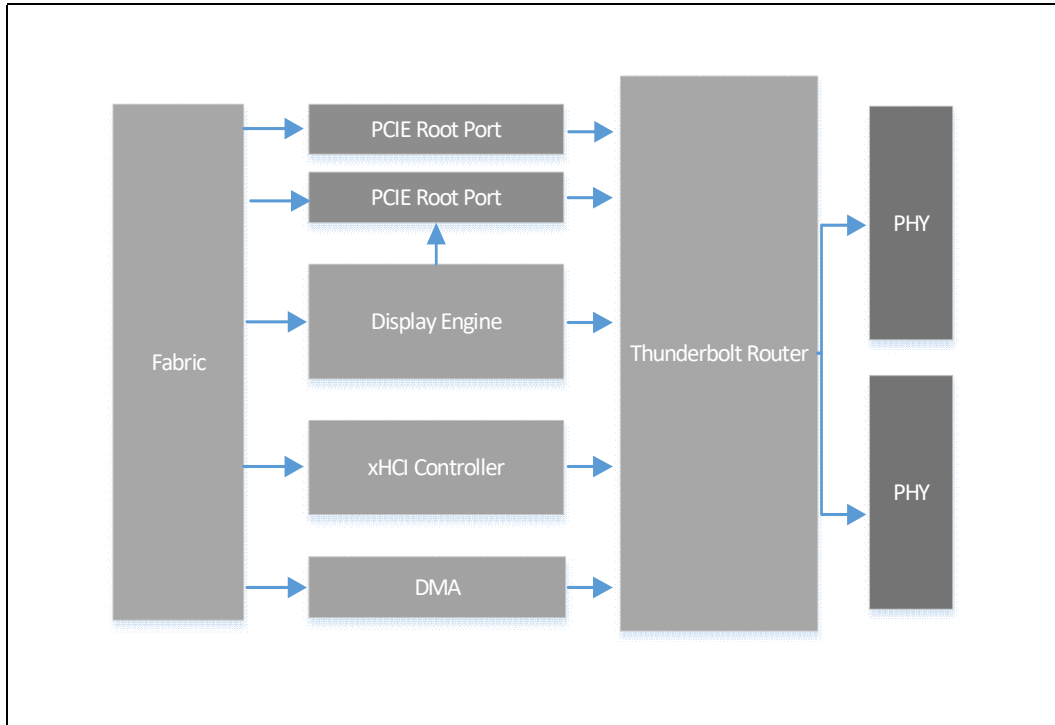
- Integrated Thunderbolt™ is a connection-oriented, tunneling architecture designed to combine multiple protocols onto a single physical interface, so that the total speed and performance of the Thunderbolt™ interface can be dynamically shared.
- The Integrated Thunderbolt™ is designed to meet the needs of multiple transport protocols and can transport native CIO packets as well as tunnel the PCI Express, DisplayPort and USB protocols.
- The integrated Thunderbolt™ controller acts as a point of entry in the CIO domain. The CIO domain is built as a daisy chain of CIO enabled products for the encapsulated protocols PCIe, DisplayPort and USB. These protocols are encapsulated into the CIO fabric and can be tunneled across the domain.
- The integrated Thunderbolt™ connection maximum data rate is 20.625 Gbps per lane but supports also 20.0 Gbps, 10.3125 Gbps and 10.0 Gbps and is compatible with older Thunderbolt™/CIO device speeds.

7.1 Host Router Implementation Capabilities

The integrated Thunderbolt™ implements the following channels.

- Two DisplayPort sink interfaces each one capable of:
 - DisplayPort 1.4 specification for tunneling
 - 1.62 Gbps or 2.7 Gbps or 5.4 Gbps or 8.1 Gbps signaling rate
 - x1, x2 or x4 lane operation
 - Support for DSC compression
- Two PCI Express Root Port interfaces each one capable of:
 - PCI Express 3.0 x4 compliant @ 8.0 GT/s
- Two xHCI Port interfaces each one capable of:
 - USB 3.1 Gen2 compliant @ 10.0 Gbps
- CIO Host Interface:
 - PCI Express 3.0 x4 compliant endpoint
 - Supports simultaneous transmit and receive on 12 paths
 - Raw mode and frame mode operation configurable on a per-path basis
 - MSI and MSI-X support
 - Interrupt moderation support
- CIO Time Management Unit (TMU):
- Two Interfaces to USB-C* connectors, each one supports:
 - Thunderbolt™ alternate mode
 - 20 paths per port
 - Each port support 20.625/20.0 Gbps or 10.3125/10.0 Gbps signaling rate
 - 16 counters per port

Figure 7-1. High Level Block Diagram



§ §



8 Graphics

8.1 Processor Graphics

The processor graphics is based on Gen11 (Generation 11) graphics core architecture that enables substantial gains in performance and lower-power consumption over prior generations. Gen 11 architecture supports up to 64 Execution Units (EUs) depending on the processor SKU.

The processor graphics architecture delivers high dynamic range of scaling to address segments spanning low power to high power, increased performance per watt, support for next generation of APIs. Gen 11 scalable architecture is partitioned by usage domains along Render/Geometry, Media, and Display. The architecture also delivers very low-power video playback and next generation analytics and filters for imaging related applications. The new Graphics Architecture includes 3D compute elements, Multi-format HW assisted decode/encode pipeline, and Mid-Level Cache (MLC) for superior high definition playback, video quality, and improved 3D performance and media.

The Display Engine handles delivering the pixels to the screen. GSA (Graphics in System Agent) is the primary channel interface for display memory accesses and "PCI-like" traffic in and out.

Table 8-1. Supported configuration by SKU

SKU	Gen	Pipes	DDI	Type-C	Thunderbolt™
Y42	Gen11 GT2	3	2	3	y
U42	Gen11 GT2	3	2	4	y

8.1.1 Media Support (Intel® QuickSync and Clear Video Technology HD)

Gen 11 implements multiple media video codecs in hardware as well as a rich set of image processing algorithms.

Note: All supported media codecs operate on 10 bpc, YCbCr 4:2:0 video profiles.

8.1.1.1 Hardware Accelerated Video Decode

Gen 11 implements a high-performance and low-power HW acceleration for video decoding operations for multiple video codecs.

The HW decode is exposed by the graphics driver using the following APIs:

- Direct3D* 9 Video API (DXVA2)
- Direct3D12 Video API Intel® Media SDK
- MFT (Media Foundation Transform) filters

Gen 11 supports full HW accelerated video decoding for AVC/VC1/MPEG2/HEVC/VP8/JPEG.



Note: HEVC – 10 bit support.

Table 8-2. Hardware Accelerated Video Decoding

Codec	Profile	Level	Maximum Resolution
MPEG2	Main	Main High	1080p
VC1/WMV9	Advanced Main Simple	L3 High Simple	3840x3840
AVC/H264	High Main	L5.2	2160p(4K)
VP8	0	Unified level	1080p
JPEG/MJPEG	Baseline	Unified level	16K x16K
HEVC/H265 (8 bits)	Main	L5.1	2160(4K)
HEVC/H265 (10 bits)	Main BT2020, isolate Dec	—	—
VP9	0 (4:2:0 Chroma 8-bit) 2 (4:2:0 Chroma 10-bit)	Unified level	2160(4K)

Expected performance:

- More than 16 simultaneous decode streams @ 1080p.

Note: Actual performance depends on the processor SKU, content bit rate, and memory frequency. Hardware decode for H264 SVC is not supported.

8.1.1.2 Hardware Accelerated Video Encode

Gen 11 implements a high-performance and low-power HW acceleration for video decoding operations for multiple video codecs.

The HW encode is exposed by the graphics driver using the following APIs:

- Intel® Media SDK
- MFT (Media Foundation Transform) filters

Gen 11 supports full HW accelerated video encoding for AVC/MPEG2/HEVC/VP9/JPEG.

Table 8-3. Hardware Accelerated Video Encode

Codec	Profile	Level	Maximum Resolution
MPEG2	Main	High	1080p
AVC/H264	High Main	L5.1	2160p(4K)
VP8	Unified profile	Unified level	—
JPEG	Baseline	—	16Kx16K
HEVC/H265	Main	L5.1	2160p(4K)
VP9	Support 8 bits 4:2:0 BT2020 may be obtained the pre/post processing	—	—



Note: Hardware encode for H264 SVC is not supported.

8.1.1.3 Hardware Accelerated Video Processing

There is hardware support for image processing functions such as De-interlacing, Film cadence detection, Advanced Video Scaler (AVS), detail enhancement, image stabilization, gamut compression, HD adaptive contrast enhancement, skin tone enhancement, total color control, Chroma de-noise, SFC pipe (Scalar and Format Conversion), memory compression, Localized Adaptive Contrast Enhancement (LACE), spatial de-noise, Out-Of-Loop De-blocking (from AVC decoder), 16 bpc support for de-noise/de-mosaic.

There is support for Hardware assisted Motion Estimation engine for AVC/MPEG2 encode, True Motion, and Image stabilization applications.

The HW video processing is exposed by the graphics driver using the following APIs:

- Direct3D* 9 Video API (DXVA2)
- Direct3D 11 Video API
- Intel® Media SDK
- MFT (Media Foundation Transform) filters
- Intel® CUI SDK

Note: Not all features are supported by all the above APIs. Refer to the relevant documentation for more details.

8.1.1.4 Hardware Accelerated Transcoding

Transcoding is a combination of decode video processing (optional) and encode. Using the above hardware capabilities can accomplish a high-performance transcode pipeline. There is not a dedicated API for transcoding.

The processor graphics supports the following transcoding features:

- Low-power and low-latency AVC encoder for video conferencing and Wireless Display applications.
- Lossless memory compression for media engine to reduce media power.
- HW assisted Advanced Video Scaler.
- Low power Scaler and Format Converter.

8.2 Platform Graphics Hardware Feature

8.2.1 Hybrid Graphics

Microsoft* Windows* 10 operating system enables the Win10 Hybrid graphics framework wherein the GPUs and their drivers can be simultaneously utilized to provide users with the benefits of both performance capability of discrete GPU (dGPU) and low-power display capability of the processor GPU (iGPU). For instance, when there is a high-end 3D gaming workload in progress, the dGPU will process and render the game frames using its graphics performance, while iGPU continues to perform the display



operations by compositing the frames rendered by dGPU. We recommend that OEMS should seek further guidance from MS to confirm that the design fits all the latest criteria defined by MS to support HG.

Microsoft* Hybrid Graphics definition includes the following:

1. The system contains a single integrated GPU and a single discrete GPU.
2. It is a design assumption that the discrete GPU has a significantly higher performance than the integrated GPU.
3. Both GPUs shall be physically enclosed as part of the system.
 - MS Hybrid DOES NOT support hot-plugging of GPUs.
 - OEMS should seek further guidance from MS before designing systems with the concept of hot-plugging.
4. Starting with Windows*10 Th1 (WDDM 2.0), a previous restriction that the discrete GPU is a render-only device, with no displays connected to it, has been removed. A render-only configuration with NO outputs is still allowed, just NOT required.

It must be noted that systems that have outputs available off of the discrete GPU will NOT support previous versions of the OS (Windows* 8.1 and Older).

Table 8-4. Hybrid Graphics Hardware Configuration

Feature	Y Processor Line	U Processor Line
PCIe* Configurations for dGFX	NA ¹	1 X 4
Hybrid Graphics	NA ¹	Yes
Note: 1. Hybrid Graphics is not POR for Y Processor Line.		





9 Display

9.1 Display Technologies Support

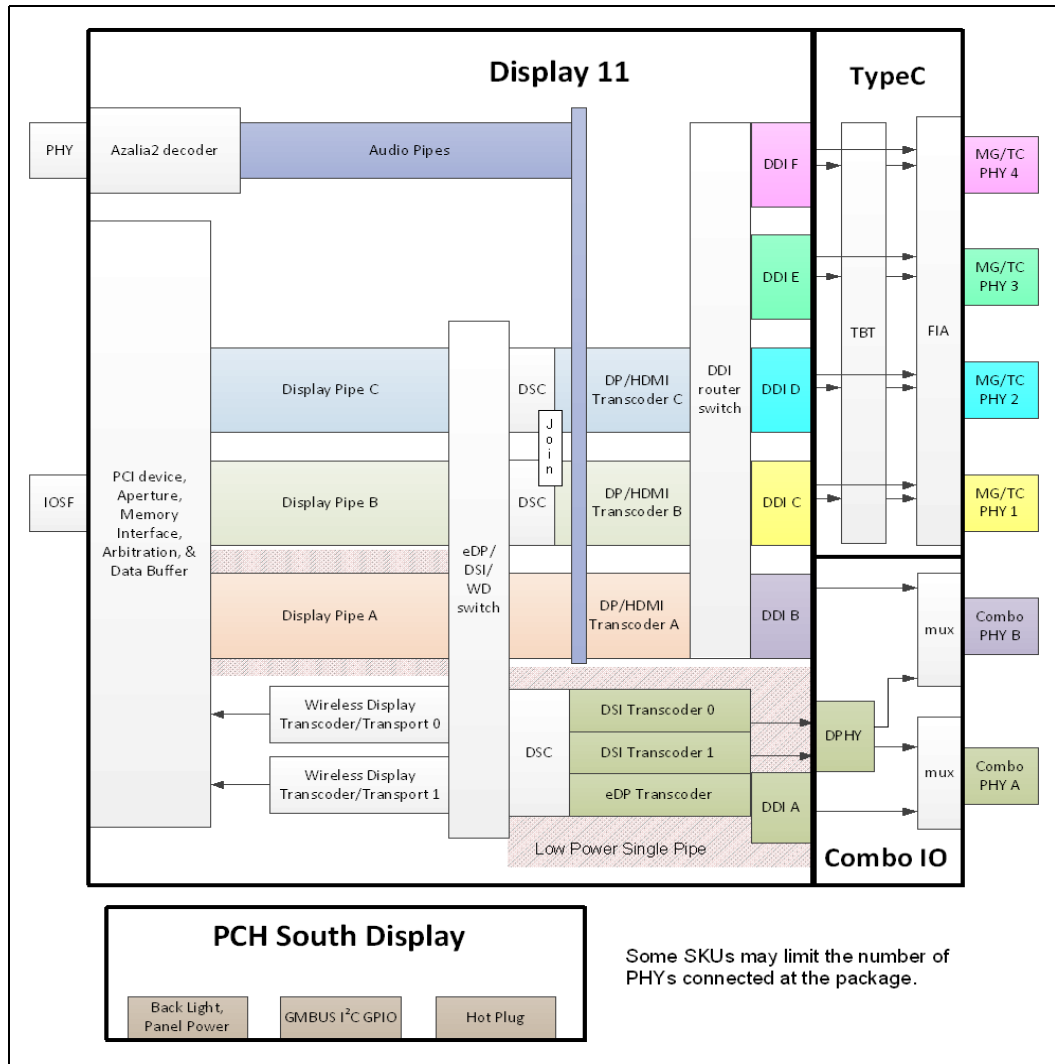
Technology	Standard
eDP* 1.4b	VESA* Embedded DisplayPort* Standard 1.4b
DisplayPort* 1.4a	VESA* DisplayPort* Standard 1.4a VESA* DisplayPort* PHY Compliance Test Specification 1.4a VESA* DisplayPort* Link Layer Compliance Test Specification 1.4a VESA* DisplayPort* Alt Mode on USB Type-C Standard Version 1.0b
HDMI* 2.0b	High-Definition Multimedia Interface Specification Version 2.0b

9.2 Display Configuration

Table 9-1. Display Ports Availability and Link Rate for Y/U-Processor Lines

SKU	Y-Processor Line 4 Core GT2	U-Processor Line 4 Core GT2
DDI A ^{1,2}	eDP* up to HBR3	eDP* up to HBR3
DDI B ²	DP* up to HBR2 HDMI* up to 5.94 Gbps	DP* up to HBR2 HDMI* up to 5.94 Gbps
USB-C* 0 (DDI C)	DP* up to HBR3 HDMI* up to 5.94 Gbps	DP* up to HBR3 HDMI* up to 5.94 Gbps
USB-C* 1 (DDI D)	DP* up to HBR3 HDMI* up to 5.94 Gbps	DP* up to HBR3 HDMI* up to 5.94 Gbps
USB-C* 2 (DDI E)	DP* up to HBR3 HDMI* up to 5.94 Gbps	DP* up to HBR3 HDMI* up to 5.94 Gbps
USB-C* 3 (DDI F)	N/A	DP* up to HBR3 HDMI* up to 5.94 Gbps
Notes:		
1. HBR3 - 8.1 Gbps lane rate.		
2. HBR2 - 5.4 Gbps lane rate.		

Figure 9-1. Processor Display Architecture



9.3 Display Features

9.3.1 General Capabilities

- Up to three simultaneous displays
- Audio stream support on external ports
- Up to four USB* Type-C for DisplayPort* Alt Mode, DisplayPort* over TBT, or DisplayPort*/HDMI* connectors
- Gamma Correction
- Color space conversion
- HDR support
- DPST - Display Power Saving Technology



- Low Power optimized pipe A
 - LACE (Localized Adaptive Contrast Enhancement), supported up to 4 K resolutions
- 3D LUT - power efficient pixel modification function for color processing

9.3.2 Multiple Display Configurations

The following multiple display configuration modes are supported (with appropriate driver software):

- Single Display is a mode with one display port activated to display the output to one display device.
- Display Clone is a mode with up to three display ports activated to drive the display content of same color depth setting but potentially different refresh rate and resolution settings to all the active display devices connected.
- Extended Desktop is a mode with up to three display ports activated to drive the content with potentially different color depth, refresh rate, and resolution settings on each of the active display devices connected.

9.3.3 High-bandwidth Digital Content Protection (HDCP)

HDCP is the technology for protecting high-definition content against unauthorized copy or unreceptive between a source (computer, digital set top boxes, and so on) and the sink (panels, monitor, and TVs). The processor supports both HDCP 2.2 and 1.4 content protection over wired displays (HDMI*, DVI, and DisplayPort*).

The HDCP 1.4/2.2 keys are integrated into the processor and customers are not required to physically configure or handle the keys.

9.3.4 DisplayPort*

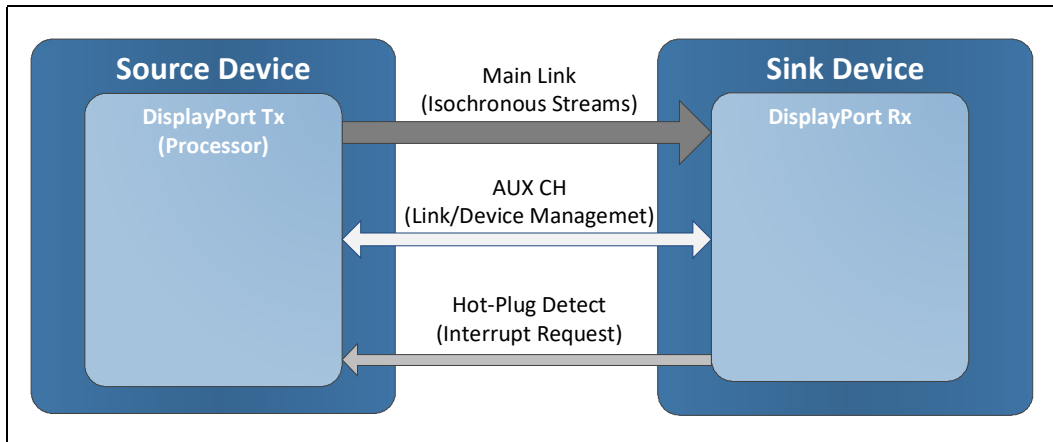
The DisplayPort* is a digital communication interface that uses differential signaling to achieve a high-bandwidth bus interface designed to support connections between PCs and monitors, projectors, and TV displays.

A DisplayPort* consists of a Main Link (4 lanes), Auxiliary channel, and a Hot-Plug Detect signal. The Main Link is a unidirectional, high-bandwidth, and low-latency channel used for transport of isochronous data streams such as uncompressed video and audio. The Auxiliary Channel (AUX CH) is a half-duplex bidirectional channel used for link management and device control. The Hot-Plug Detect (HPD) signal serves as an interrupt request from the sink device to source device.

The processor is designed in accordance to VESA* DisplayPort* specification. Refer to [Section 9.1, "Display Technologies Support"](#).

The DisplayPort* support DisplayPort* Alt mode over Type-C and DP tunneling via TBT. Refer to [Chapter 6, "USB-C* Sub System"](#) For DisplayPort* Alt mode support and [Chapter 7, "Thunderbolt™"](#) for DisplayPort* tunneling.

Figure 9-2. DisplayPort* Overview



- Support main link of 1, 2, or 4 data lanes.
- Aux channel for Link/Device management.
- Support up to 36 BPP (Bit Per Pixel).
- Support SSC.
- Support YCbCR 4:4:4, YCbCR 4:2:0, and RGB color format.
- Support MST (Multi-Stream Transport).
- Support VESA DSC 1.1.
- Adaptive sync.

9.3.4.1 Multi-Stream Transport (MST)

- The processor supports Multi-Stream Transport (MST), enabling multiple monitors to be used via a single DisplayPort connector.
- MST does not supported concurrent with DSC.
- Max MST DP supported resolution.

Table 9-2. Display Resolutions and Link Bandwidth for Multi-Stream Transport Calculations (Sheet 1 of 2)

Pixels per line	Lines	Refresh Rate [Hz]	Pixel Clock [MHz]	Link Bandwidth [Gbps]
640	480	60	25.2	0.76
800	600	60	40	1.20
1024	768	60	65	1.95
1280	720	60	74.25	2.23
1280	768	60	68.25	2.05
1360	768	60	85.5	2.57
1280	1024	60	108	3.24
1400	1050	60	101	3.03
1680	1050	60	119	3.57
1920	1080	60	148.5	4.46



Table 9-2. Display Resolutions and Link Bandwidth for Multi-Stream Transport Calculations (Sheet 2 of 2)

Pixels per line	Lines	Refresh Rate [Hz]	Pixel Clock [MHz]	Link Bandwidth [Gbps]
1920	1200	60	154	4.62
2048	1152	60	156.75	4.70
2048	1280	60	174.25	5.23
2048	1536	60	209.25	6.28
2304	1440	60	218.75	6.56
2560	1440	60	241.5	7.25
3840	2160	30	262.75	7.88
2560	1600	60	268.5	8.06
2880	1800	60	337.5	10.13
3200	2400	60	497.75	14.93
3840	2160	60	533.25	16.00
4096	2160	60	556.75	16.70
4096	2304	60	605	18.15
5120	3200	60	1042.5	31.28

Notes:

- All above is related to bit depth of 24.
- The data rate for a given video mode can be calculated as: Data Rate = Pixel Frequency * Bit Depth.
- The bandwidth requirements for a given video mode can be calculated as: Bandwidth = Data Rate * 1.25 (for 8B/10B coding overhead).
- The link bandwidth depends if the standards is reduced blanking or not.
If the standard is not reduced blanking - the expected bandwidth may be higher.
For more details refer to VESA and Industry Standards and Guidelines for Computer Display Monitor Timing (DMT).Version 1.0, Rev. 13 February 8, 2013.
- To calculate what are the resolutions that can be supported in MST configurations, follow the below guidelines:
 - Identify what is the link bandwidth column according to the requested display resolution.
 - Summarize the bandwidth for two of three displays accordingly, and make sure the final result is below 21.6Gbps. (for example: 4 lanes HBR2 bit rate).
 For example:
 - Docking two displays: 3840x2160@60 Hz + 1920x1200@60 Hz = 16 + 4.62 = 20.62 Gbps [Supported].
 - Docking three displays: 3840x2160@30 Hz + 3840x2160@30 Hz + 1920x1080@60 Hz = 7.88 + 7.88 + 4.16 = 19.92 Gbps [Supported].

Table 9-3. DisplayPort* Maximum Resolution

Standard	Y-Processor Line ¹	U-Processor Line ¹
DP*	4096x2304 60Hz 36 bpp 5120x3200 60Hz 24 bpp	4096x2304 60Hz 36 bpp 5120x3200 60Hz 24 bpp
DP* with DSC	5120x3200 60Hz 36 bpp	5120x3200 60Hz 36 bpp

Notes:

- Maximum resolution is based on implementation of 4 lanes at HBR3 link data rate.
- bpp - bit per pixel.
- Resolution support are subject to memory BW availability.

9.3.5 High-Definition Multimedia Interface (HDMI*)

The High-Definition Multimedia Interface (HDMI*) is provided for transmitting uncompressed digital audio and video signals from DVD players, set-top boxes, and other audio-visual sources to television sets, projectors, and other video displays. It can carry high-quality multi-channel audio data and all standard and high-definition

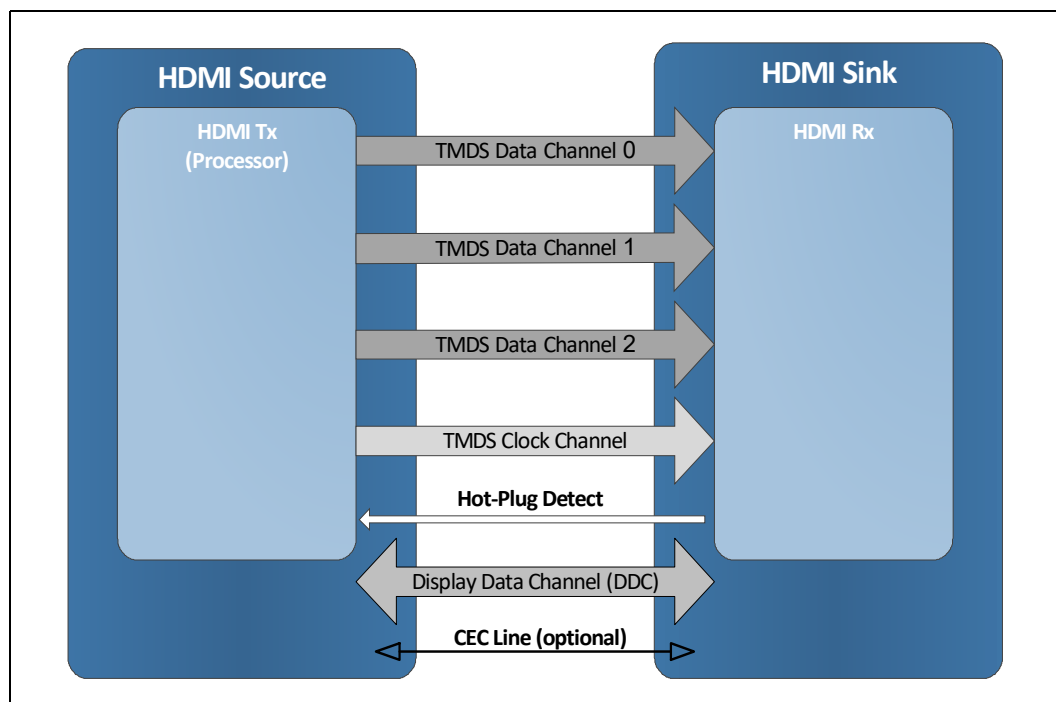
consumer electronics video formats. The HDMI display interface connecting the processor and display devices uses transition minimized differential signaling (TMDS) to carry audiovisual information through the same HDMI cable.

HDMI includes three separate communications channels: TMDS, DDC, and the optional CEC (consumer electronics control). CEC is not supported on the processor. As shown in the following figure, the HDMI cable carries four differential pairs that make up the TMDS data and clock channels. These channels are used to carry video, audio, and auxiliary data. In addition, HDMI carries a VESA DDC. The DDC is used by an HDMI Source to determine the capabilities and characteristics of the Sink.

Audio, video, and auxiliary (control/status) data is transmitted across the three TMDS data channels. The video pixel clock is transmitted on the TMDS clock channel and is used by the receiver for data recovery on the three data channels. The digital display data signals driven natively through the PCH are AC coupled and needs level shifting to convert the AC coupled signals to the HDMI compliant digital signals.

The processor HDMI interface is designed in accordance with the High-Definition Multimedia Interface.

Figure 9-3. HDMI* Overview



- DDC (Display Data Channel) channel.
- Support YCbCR 4:4:4, YCbCR 4:2:0, and RGB color format.
- Support up to 36 BPP (Bit Per Pixel).

**Table 9-4. HDMI* Maximum Resolution**

Standard	Y-Processor Line ¹	U-Processor Line ¹
HDMI 1.4	4Kx2K 24-30Hz 24 bpp	4Kx2K 24-30 Hz 24 bpp
HDMI 2.0b	4Kx2K 48-60 Hz 24 bpp (RGB/ YUV444) 4Kx2K 48-60 Hz 12 bpc (YUV420)	4Kx2K 48-60 Hz 24 bpp (RGB/ YUV444) 4Kx2K 48-60 Hz 12 bpc (YUV420)
Notes: 1. bpp - bit per pixel. 2. Resolution support are subject to memory BW availability.		

9.3.6 Digital Video Interface (DVI)

The processor Digital Ports can be configured to drive DVI-D. DVI uses TMDS for transmitting data from the transmitter to the receiver, which is similar to the HDMI protocol except for the audio and CEC. Refer to the HDMI section for more information on the signals and data transmission. The digital display data signals driven natively through the processor are AC coupled and need level shifting to convert the AC coupled signals to the HDMI compliant digital signals.

Table 9-5. DVI Maximum Resolution Supported

Standard	Y-Processor Line	U-Processor Line
DVI	1920x1200 60 Hz 24 bpp	1920x1200 60 Hz 24 bpp
Notes: 1. bpp - bit per pixel. 2. Resolution support are subject to memory BW availability.		

9.3.7 embedded DisplayPort* (eDP*)

The embedded DisplayPort* (eDP*) is an embedded version of the DisplayPort* standard oriented towards applications such as notebook and All-In-One PCs. Like DisplayPort*, embedded DisplayPort* also consists of a Main Link, Auxiliary channel, and an optional Hot-Plug Detect signal.

- Supported on Low power optimized pipe A
- Support up to HBR3 link rate
- Support Backlight PWM control signal
- Support VESA DSC (Data Stream Compression)
- Support SSC
- Panel Self Refresh 1
- Panel Self Refresh 2
- MSO 2x2 (Multi Segment Operation)
- Dedicated Aux channel
- Adaptive sync

Table 9-6. Embedded DisplayPort Maximum Resolution (Sheet 1 of 2)

Standard	Y-Processor Line ¹	U-Processor Line ¹
eDP*	4096x2304 60 Hz 36 bpp 5120x3200 60 Hz 24 bpp	4096x2304 60 Hz 36 bpp 5120x3200 60 Hz 24 bpp
eDP* with DSC	5120x3200 60 Hz 36bpp	5120x3200 60 Hz 36 bpp



Table 9-6. Embedded DisplayPort Maximum Resolution (Sheet 2 of 2)

Standard	Y-Processor Line ¹	U-Processor Line ¹
Notes: 1. Maximum resolution is based on implementation of 4 lanes at HBR3 link data rate. 2. PSR2 supported for up to 4K resolutions. 3. bpp - bit per pixel. 4. Resolution support are subject to memory BW availability.		

9.3.8 Integrated Audio

- HDMI* and DisplayPort interfaces carry audio along with video.
- The processor supports three High Definition audio streams on three digital ports simultaneously (the DMA controllers are in PCH).
- The integrated audio processing (DSP) is performed by the PCH, and delivered to the processor using the AUDIO_SDI and AUDIO_CLK inputs pins.
- AUDIO_SDO output pin is used to carry responses back to the PCH.
- Supports only the internal HDMI and DP CODECs.

Table 9-7. Processor Supported Audio Formats over HDMI and DisplayPort*

Audio Formats	HDMI*	DisplayPort*
AC-3 Dolby* Digital	Yes	Yes
Dolby Digital Plus	Yes	Yes
DTS-HD*	Yes	Yes
LPCM, 192 kHz/24 bit, 6 Channel	Yes	Yes
Dolby TrueHD, DTS-HD Master Audio* (Lossless Blu-Ray Disc* Audio Format)	Yes	Yes

The processor will continue to support Silent stream. Silent stream is an integrated audio feature that enables short audio streams, such as system events to be heard over the HDMI* and DisplayPort* monitors. The processor supports silent streams over the HDMI and DisplayPort interfaces at 48 kHz sample-rate two channel support.

Note: 88.2 kHz, 96 kHz, 176.4 kHz, and 192 kHz sampling rates and multi-channel silent stream support are being evaluated.





10 Camera/MIPI

10.1 Camera Pipe Support

Camera pipe functions such as de-mosaic, white balance, defect pixel correction, black level correction, gamma correction, vignette control, Front end Color Space Converter (CSC), Image Enhancement Color Processing (IECP).

10.2 MIPI* CSI-2 Camera Interconnect

The Camera I/O Controller provides a native/integrated interconnect to camera sensors, compliant with MIPI DPHY1.2 CSI2 V1.3 protocol. A total of 32 (U Processor Line) and 38 (Y Processor Line) lanes are available for the camera interface supporting up to 6 sensors at U Processor segment and up to 7 sensors at Y Processor segment.

Data transmission interface (referred as CSI-2) is a unidirectional differential serial interface with data and clock signals; the physical layer of this interface is the MIPI* Alliance Specification for D-PHY.

The control interface (referred as CCI) is a bi-directional control interface compatible with I²C standard.

Note: The CSI-2 interface is available only on Y-Processor Line and U-Processor Line.

10.2.1 Camera Control Logic

The camera infrastructure supports several architectural options for camera control utilizing camera PMIC and/or discrete logic. IPU4 control options utilize I²C for bidirectional communication and PCH GPIOs to drive various control functions.

10.2.2 Camera Modules

Intel maintains an Intel User Facing Camera Approved Vendor List and Intel World-Facing Approved Vendor List to simplify system design. Additional services are available to support non-AVL options.



10.2.3 CSI-2 Lane Configuration

Port Data/Clock	Configuration Option 1	Port Data/Clock	Configuration Option 2
Port D Clock	x4	Port D Clock	x2
Port D Lane 0		Port D Lane 0	
Port D Lane 1		Port D Lane 1	
Port D Lane 2		Port C Lane 0	x1
Port D Lane 3		Port C Clock	
Port E Clock	x2		
Port E Lane 0			
Port E Lane 1			
Port F Clock	x2		
Port F Lane 0			
Port F Lane 1			
Port H Clock	x4	Port H Clock	x2
Port H Lane 0		Port H Lane 0	
Port H Lane 1		Port H Lane 1	
Port H Lane 2		Port G Lane 0	x1
Port H Lane 3		Port G Clock	
Port A ² Lane 0	x2		
Port A ² Lane 1			
Port A ² Lane 2			

Notes:

1. In Configuration Option 1 the pin is functioning as Port D (**Data**) Lane 3 while in Configuration option 2 the pin is functioning as Port C **Clock**, the same apply to Port H Lane 3 **DATA** and Port G **Clock**.
2. Port A Available in Y Processor Line only.
3. All lanes are DPHY1.2 up to 2.5 Gbps.

For implementation and more information, contact the Intel representative.





11 Signal Description

This chapter describes the processor signals. They are arranged in functional groups according to their associated interface or category. The notations in the following table are used to describe the signal type.

The signal description also includes the type of buffer used for the particular signal (refer the following table).

Table 11-1. Signal Tables Terminology

Notation	Signal Type
I	Input pin
O	Output pin
I/O	Bi-directional Input/Output pin
SE	Single Ended Link
Diff	Differential Link
CMOS	CMOS buffers. 1.05V- tolerant
OD	Open Drain buffer
LPDDR4/x	LPDDR4/LPDDR4x buffers: 1.1V-tolerant
DDR4	DDR4 buffers: 1.2V-tolerant
A	Analog reference or output. May be used as a threshold voltage or for buffer compensation
GTL	Gunning Transceiver Logic signaling technology
Ref	Voltage reference signal
Availability	Signal Availability condition - based on segment, SKU, platform type or any other factor
Asynchronous ¹	Signal has no timing relationship with any reference clock.
Note: 1. Qualifier for a buffer type.	

11.1 System Memory Interface

11.1.1 DDR4 Memory Interface

Table 11-2. DDR4 Memory Interface (Sheet 1 of 3)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDR0_DQ[7:0][[7:0]] DDR1_DQ[7:0][[7:0]]	Data Buses: Data signals interface to the SDRAM data buses. Example: DDR0_DQ2[5]	I/O	DDR4	SE	U-Processor Line
DDR0_DQSP[7:0] DDR0_DQSN[7:0] DDR1_DQSP[7:0] DDR1_DQSN[7:0]	Data Strobes: Differential data strobe pairs. The data is captured at the crossing point of DQS during read and write transactions.	I/O	DDR4	Diff	U-Processor Lines



Table 11-2. DDR4 Memory Interface (Sheet 2 of 3)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDR0_CLK_N[1:0] DDR0_CLK_P[1:0] DDR1_CLK_N[1:0] DDR1_CLK_P[1:0]	SDRAM Differential Clock: Differential clocks signal pairs, pair per rank. The crossing of the positive edge of DDR0_CLK_P/DDR1_CLK_P and the negative edge of their complement DDR0_CLK_N / DDR1_CLK_N are used to sample the command and control signals on the SDRAM.	O	DDR4	Diff	U-Processor Line
DDR0_CKE[1:0] DDR1_CKE[1:0]	Clock Enable: (1 per rank). These signals are used to: <ul style="list-style-type: none"> Initialize the SDRAMs during power-up. Power-down SDRAM ranks. Place all SDRAM ranks into and out of self-refresh during STR (Suspend to RAM). 	O	DDR4	SE	U-Processor Line,
DDR0_CS#[1:0] DDR1_CS#[1:0]	Chip Select: (1 per rank). These signals are used to select particular SDRAM components during the active state. There is one Chip Select for each SDRAM rank.	O	DDR4	SE	U-Processor Line
DDR0_ODT[1:0] DDR1_ODT[1:0]	On Die Termination: (1 per rank). Active SDRAM Termination Control.	O	DDR4	SE	U-Processor Line
DDR0_MA[16:0] DDR1_MA[16:0]	Address: These signals are used to provide the multiplexed row and column address to the SDRAM. DDR0_MA[16] uses as RAS# signal DDR0_MA[15] uses as CAS# signal DDR0_MA[14] uses as WE# signal DDR1_MA[16] uses as RAS# signal DDR1_MA[15] uses as CAS# signal DDR1_MA[14] uses as WE# signal	O	DDR4	SE	U-Processor Line
DDR0_ACT# DDR1_ACT#	Activation Command: ACT# HIGH along with CS_N determines that the signals addresses below have command functionality.	O	DDR4	SE	U-Processor Line
DDR0_BG[1:0] DDR1_BG[1:0]	Bank Group: BG[0:1] define to which bank group an Active, Read, Write or Precharge command is being applied. BG0 also determines which mode register is to be accessed during a MRS cycle.	O	DDR4	SE	U-Processor Line For DDP, BG[1] should be connected.
DDR0_BA[1:0] DDR1_BA[1:0]	Bank Address: BA[1:0] define to which bank an Active, Read, Write or Precharge command is being applied. Bank address also determines which mode register is to be accessed during a MRS cycle.	O	DDR4	SE	U-Processor Line
DDR0_ALERT# DDR1_ALERT#	Alert: This signal is used at command training only. It is getting the Command and Address Parity error flag during training. CRC feature is not supported.	I	DDR4	SE	U-Processor Line



Table 11-2. DDR4 Memory Interface (Sheet 3 of 3)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDR0_PAR DDR1_PAR	Command and Address Parity: These signals are used for parity check.	O	A	SE	U-Processor Line
DDR0_VREF_CA DDR1_VREF_CA	Memory Reference Voltage for Command and Address: Refer to the appropriate design guidelines for implementation details.	O	A	SE	U-Processor Line
DDR_RCOMP[2:0]	System Memory Resistance Compensation: Refer to the appropriate design guidelines for implementation details and values.	N/A	A	SE	U-Processor Line
DRAM_RESET#	Memory Reset: Refer to the appropriate design guidelines for implementation details.	O	CMOS	SE	U-Processor Line
DDR_VTT_CTL	System Memory Power Gate Control: When signal is high – platform memory VTT regulator is enable, output high. When signal is low - Disables the platform memory VTT regulator in C8 and deeper and S3.	O	A	SE	U-Processor Line

11.1.2 LPDDR4 Memory Interface

Table 11-3. LPDDR4 Memory Interface (Sheet 1 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDRA_DQ[3:0][7:0] DDRB_DQ[3:0][7:0] DDRC_DQ[3:0][7:0] DDRD_DQ[3:0][7:0]	Data Buses: Data signals interface to the SDRAM data buses. Example: DDR0_DQ2[5]	I/O	LPDDR4	SE	U/Y Processor Line
DDRA_DQSP[3:0] DDRB_DQSP[3:0] DDRC_DQSP[3:0] DDRD_DQSP[3:0] DDRA_DQSN[3:0] DDRB_DQSN[3:0] DDRC_DQSN[3:0] DDRD_DQSN[3:0]	Data Strobes: Differential data strobe pairs. The data is captured at the crossing point of DQS during read and write transactions.	I/O	LPDDR4	Diff	U/Y Processor Line
DDRA_CLK_N DDRA_CLK_P DDRB_CLK_N DDRB_CLK_P DDRC_CLK_N DDRC_CLK_P DDRD_CLK_N DDRD_CLK_P	SDRAM Differential Clock: Differential clocks signal pairs, pair per channel and package. The crossing of the positive edge of DDRA_CLK_P, DDRB_CLK_P, DDRC_CLK_P, DDRD_CLK_P and the negative edge of their complement DDRA_CLK_N, DDRB_CLK_N, DDRC_CLK_N, DDRD_CLK_N are used to sample the command and control signals on the SDRAM.	I/O	LPDDR4	Diff	U/Y Processor Line

Table 11-3. LPDDR4 Memory Interface (Sheet 2 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDRA_CKE[1:0] DDRB_CKE[1:0] DDRC_CKE[1:0] DDRD_CKE[1:0]	Clock Enable: (1 per rank) These signals are used to: <ul style="list-style-type: none"> Initialize the SDRAMs during power-up. Power-down SDRAM ranks. Place all SDRAM ranks into and out of self-refresh during STR. 	O	LPDDR4	SE	U/Y Processor Line
DDRA_CS[1:0] DDRB_CS[1:0] DDRC_CS[1:0] DDRD_CS[1:0]	Chip Select: (1 per rank). These signals are used to select particular SDRAM components during the active state. There is one Chip Select for each SDRAM rank. The Chip select signal is Active High.	O	LPDDR4	SE	U/Y Processor Line
DDRA_CA[5:0] DDRB_CA[5:0] DDRC_CA[5:0] DDRD_CA[5:0]	Command Address: These signals are used to provide the multiplexed command and address to the SDRAM.	O	LPDDR4	SE	U/Y Processor Line
DDR_RCOMP[2:0]	System Memory Resistance Compensation: Refer to the appropriate design guidelines for implementation details and values.	O	A	SE	U/Y Processor Line
DRAM_RESET#	Memory Reset: Refer to the appropriate design guidelines for implementation details.	O	CMOS	SE	U/Y Processor Line

11.2 Reset and Miscellaneous Signals

Table 11-4. Reset and Miscellaneous Signals (Sheet 1 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
CFG[19:0]	Configuration Signals: The CFG signals have a default value of '1' if not terminated on the board. Refer to the appropriate design guidelines for pull-down recommendations when a logic low is desired. Intel recommends placing test points on the board for CFG pins. <ul style="list-style-type: none"> CFG[0]: Stall reset sequence after PCU PLL lock until de-asserted: <ul style="list-style-type: none"> 1 = (Default) Normal Operation; No stall. 0 = Stall. CFG[3:1]: Reserved configuration lane. CFG[4]: eDP enable: <ul style="list-style-type: none"> 1 = Disabled. 0 = Enabled. CFG[19:5]: Reserved configuration lanes. 	I	GTL	SE	U/Y Processor Lines
CFG_RCOMP	Configuration Resistance Compensation	N/A	N/A	SE	U/Y Processor Lines



Table 11-4. Reset and Miscellaneous Signals (Sheet 2 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
PROC_POPIRCOMP	POPIO Resistance Compensation	N/A	N/A	SE	U/Y Processor Lines
PROC_SELECT#	Processor Select: This pin is for compatibility with future platforms. It should be unconnected for 10 th Generation Intel® Core™ processor.			N/A	U Processor Lines

11.3 Display Interfaces

11.3.1 Embedded DisplayPort* (eDP*) Signals

Table 11-5. embedded DisplayPort* Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDIA_TXP[3:0] DDIA_TXN[3:0]	embedded DisplayPort Transmit: Differential pair.	O	eDP	Diff	All Processor Lines
DDIA_AUXP DDIA_AUXN	embedded DisplayPort Auxiliary: Half-duplex, bidirectional channel consist of one differential pair.	O	eDP	Diff	All Processor Lines
DISP_UTILS	embedded DisplayPort Utility: Output control signal used for brightness correction of embedded LCD displays with backlight modulation. This pin will co-exist with functionality similar to existing BKLCTCL pin on PCH.	O	Async CMOS	SE	All Processor Lines
DP_RCOMP	DDI IO Compensation resistor, supporting DP*, eDP* and HDMI* channels.	N/A	A	SE	All Processor Lines

Note: eDP* implementation go along with additional sideband signals

11.3.2 Digital Display Interface (DDI) Signals

Table 11-6. Display Interface Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDIA_TXP[3:0] DDIA_TXN[3:0] DDIB_TXP[3:0] DDIB_TXN[3:0]	Digital Display Interface Transmit: Differential Pairs.	O	Combo I/O	Diff	U/Y Processor Lines.
DDIA_AUXP DDIA_AUXN DDIB_AUXP DDIB_AUXN	Digital Display Interface Display Port Auxiliary: Half-duplex, bidirectional channel consist of one differential pair for each channel.	O	Combo I/O	Diff	



11.4 USB Type-C Signals

Table 11-7. USB Type-C Signals

Signal Name	Description	Dir.	Link Type	Availability
TCP[2:0]_TX_P[1:0] TCP[2:0]_TX_N[1:0]	TX Data Lane.	O	Diff	U/Y Processor Lines
TCP[3]_TX_P[1:0] TCP[3]_TX_N[1:0]	TX Data Lane.	O	Diff	U Processor Lines
TCP[2:0]_TXRX_P[1:0] TCP[2:0]_TXRX_N[1:0]	RX Data Lane, also serves as the secondary TX data lane.	I/O	Diff	U/Y Processor Lines
TCP[3]_TXRX_P[1:0] TCP[3]_TXRX_N[1:0]	RX Data Lane, also serves as the secondary TX data lane.	I/O	Diff	U Processor Lines
TCP[2:0]_AUXPAD_P TCP[2:0]_AUXPAD_N	Common Lane AUX-PAD.	I/O	Diff	U/Y Processor Lines
TCP[3]_AUXPAD_P TCP[3]_AUXPAD_N	Common Lane AUX-PAD.	I/O	Diff	U Processor Lines
TC_RCOMP_P TC_RCOMP_N	Type C Resistance Compensation.	N/A	Diff	U/Y Processor Lines

11.5 MIPI* CSI-2 Interface Signals

Table 11-8. MIPI* CSI-2 Interface Signals (Sheet 1 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
CSI_A_DP[0] CSI_A_DN[0]	CSI-2 Ports A Data lane	I	DPHY	Diff	Y Processor Line
CSI_C_DP[0] CSI_C_DN[0]					U/Y Processor Lines
CSI_D_DP[3:0] CSI_D_DN[3:0]	U/Y Processor Lines				
CSI_E_DP[1:0] CSI_E_DN[1:0]	U/Y Processor Lines				
CSI_F_DP[1:0] CSI_F_DN[1:0]	U/Y Processor Lines				
CSI_G_DP[0] CSI_G_DN[0]	U/Y Processor Lines				
CSI_H_DP[3:0] CSI_H_DN[3:0]	U/Y Processor Lines				
	CSI-2 Ports C-H Data lanes				



Table 11-8. MIPI* CSI-2 Interface Signals (Sheet 2 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
CSI_A_CLK_P CSI_A_CLK_N	CSI-2 Ports A Clock lane	I	DPHY	Diff	Y Processor Line
CSI_C_CLK_P CSI_C_CLK_N					U/Y Processor Lines
CSI_D_CLK_P CSI_D_CLK_N					U/Y Processor Lines
CSI_E_CLK_P CSI_E_CLK_N					U/Y Processor Lines
CSI_F_CLK_P CSI_F_CLK_N					U/Y Processor Lines
CSI_G_CLK_P CSI_G_CLK_N					U/Y Processor Lines
CSI_H_CLK_P CSI_H_CLK_N					U/Y Processor Lines
CSI_RCOMP					CSI Resistance Compensation

11.6 Testability Signals

Table 11-9. Testability Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
BPM#[3:0]	Breakpoint and Performance Monitor Signals: Outputs from the processor that indicate the status of breakpoints and programmable counters used for monitoring processor performance.	I/O	GTL	SE	U/Y Processor Lines
PROC_PRDY#	Probe Mode Ready: PROC_PRDY# is a processor output used by debug tools to determine processor debug readiness.	O	OD	SE	U/Y Processor Lines
PROC_PREQ#	Probe Mode Request: PROC_PREQ# is used by debug tools to request debug operation of the processor.	I	GTL	SE	U/Y Processor Lines
PROC_TCK	Test Clock: This signal provides the clock input for the processor Test Bus (also known as the Test Access Port). This signal should be driven low or allowed to float during power on Reset.	I	GTL	SE	U/Y Processor Lines
PROC_TDI	Test Data In: This signal transfers serial test data into the processor. This signal provides the serial input needed for JTAG specification support.	I	GTL	SE	U/Y Processor Lines
PROC_TDO	Test Data Out: This signal transfers serial test data out of the processor. This signal provides the serial output needed for JTAG specification support.	O	OD	SE	U/Y Processor Lines
PROC_TMS	Test Mode Select: A JTAG specification support signal used by debug tools.	I	GTL	SE	U/Y Processor Lines
PROC_TRST#	Test Reset: Resets the Test Access Port (TAP) logic. This signal should be driven low during power on Reset.	I	GTL	SE	U/Y Processor Lines



11.7 Error and Thermal Protection Signals

Table 11-10. Error and Thermal Protection Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
CATERR#	Catastrophic Error: This signal indicates that the system has experienced a catastrophic error and cannot continue to operate. The processor will set this signal for non-recoverable machine check errors or other unrecoverable internal errors. CATERR# is used for signaling the following types of errors: Legacy MCERRs, CATERR# is asserted for 16 BCLKs. Legacy IERRs, CATERR# remains asserted until warm or cold reset.	O	OD	SE	All Processor Lines
PECI	Platform Environment Control Interface: A serial sideband interface to the processor. It is used primarily for thermal, power, and error management. Details regarding the PEFI electrical specifications, protocols and functions can be found in the RS-Platform Environment Control Interface (PECI) Specification, Revision 3.0.	I/O	PECI, Async	SE	All Processor Lines
PROCHOT#	Processor Hot: PROCHOT# goes active when the processor temperature monitoring sensor(s) detects that the processor has reached its maximum safe operating temperature. This indicates that the processor Thermal Control Circuit (TCC) has been activated, if enabled. This signal can also be driven to the processor to activate the TCC.	I/O	GTL I OD O	SE	All Processor Lines
THRMTRIP#	Thermal Trip: The processor protects itself from catastrophic overheating by use of an internal thermal sensor. This sensor is set well above the normal operating temperature to ensure that there are no false trips. The processor will stop all executions when the junction temperature exceeds approximately 125 °C. This is signaled to the system by the THRMTRIP# pin. Refer to the appropriate design guidelines for termination requirements.	O	OD	SE	All Processor Lines

11.8 Power Sequencing Signals

Table 11-11. Power Sequencing Signals (Sheet 1 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
PROCPWRGD	Processor Power Good: The processor requires this input signal to be a clean indication that the V _{CC} and V _{DDQ} power supplies are stable and within specifications. This requirement applies regardless of the S-state of the processor. 'Clean' implies that the signal will remain low (capable of sinking leakage current), without glitches, from the time that the power supplies are turned on until they come within specification. The signal should then transition monotonically to a high state.	I	CMOS	SE	U/Y Processor Lines



Table 11-11. Power Sequencing Signals (Sheet 2 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
VCCST_OVERRIDE	VCCST_OVERRIDE: output signal from the PCH to keep VCCST powered ON (in case VCCST is powered down) for Type C wake capability (Connected to VCCST_PWRGD_TCSS on board).	O	N/A	N/A	U/Y Processor Lines
VCCST_PWRGOOD	VCCST Power Good: The processor requires this input signal to be a clean indication that the VCCST and VDDQ power supplies are stable and within specifications. This signal should have a valid level during both S0 and S3 power states. 'Clean' implies that the signal will remain low (capable of sinking leakage current), without glitches, from the time that the power supplies are turned on until they come within specification. The signal then transition monotonically to a high state.	I	CMOS	SE	U/Y Processor Lines
VCCST_PWRGD_TCSS	VCCST_PWRGD_TCSS: The processor requires this input signal to be asserted when the type-c subsystem requires keeping VCCST supply on (VCCST_OVERRIDE), even when entering S3 – S5 states. This signal start as low and may change polarity only at the entry to S3 – S5. If required to toggle, the signal level must always change before the de-assertion of VCCST_PWRGD signal at the Sx entry flow. This signal must have a valid level during S0 – S5 power states.	I	CMOS	SE	U/Y Processor Lines
SKTOCC#	Socket Occupied: Pulled down directly (0 Ohms) on the processor package to the ground. There is no connection to the processor silicon for this signal. System board designers may use this signal to determine if the processor is present.	N/A	N/A	SE	U/Y Processor Lines
VIDSOUT	VIDSOUT, VIDSCK, VIDALERT#: These signals comprise a three-signal serial synchronous interface used to transfer power management information between the processor and the voltage regulator controllers.	I/O	I:GTL/ O:OD	SE	U/Y Processor Lines
VIDSCK		O	OD		
VIDALERT#		I	CMOS		

11.9 Processor Power Rails

Table 11-12. Processor Power Rails Signals (Sheet 1 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
VCC _{IN}	On-Package VR (OPVR) power rail	I	Power	-	U/Y-Processor Line
VCC _{IN_AUX}	On-Package VR (OPVR) power auxiliary rail	I	Power	-	U/Y-Processor Line
VCC _{1p8A}	System Agent Power Rail	I	Power	-	U/Y-Processor Line
V _{DDQ}	System Memory power rail	I	Power	-	U/Y-Processor Line

Table 11-12. Processor Power Rails Signals (Sheet 2 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
VCC _{ST}	Sustain voltage for processor standby modes	I	Power	-	U/Y-Processor Line
VCC _{STG}	Gated sustain voltage for processor standby modes	I	Power	-	U/Y-Processor Line
VCC _{PLL}	Processor PLLs power rails	I	Power	-	U/Y-Processor Line
VCC _{PLL_OC}	Processor PLLs power rails	I	Power	-	U/Y-Processor Line
VCC _{IN_SENSE} VCC _{IN_AUX_VCCSENSE}	Isolated, low impedance voltage sense pins. They can be used to sense or measure voltage near the silicon.	N/A	PWR_SENSE	-	U/Y-Processor Line
VCC _{IN_AUX_VSSSENSE} VSS _{IN_SENSE}	Isolated, low impedance reference ground sense pins. They can be used to sense or measure the reference ground to the adequate voltage rail near the silicon.		GND_SENSE		U/Y-Processor Line

Table 11-13. Processor Pull-up Power Rails Signals

Signal Name	Description	Dir.	Type	Availability
VCC _{STG_OUT_LGC}	Reference power rail for all Legacy Signals Pull-up on platform.	O	Reference Power	U-Processor Line
VCC _{ST_OUT}	Reference power rail for Legacy Signals Pull-up on platform.	O	Reference Power	Y-Processor Line
VCC _{STG_OUT}	Reference power rail for JTAG/PROCHOT Signals Pull-up on platform, Supplier of the FPGM power rail.	O	Reference Power	Y-Processor Line
	VCC _{STG_OUT} Power rail.	O	Power	U-Processor Line
VCC _{IO_OUT}	Reference power rail for all Debug/Config Signals Pull-up on platform.	O	Reference Power	U/Y-Processor Line

11.10 Ground, Reserved and Non-Critical to Function (NCTF) Signals

The following are the general types of reserved (RSVD) signals and connection guidelines:

- RSVD – these signals should not be connected.
- RSVD_TP – these signals should be routed to a test point.
- _NCTF – these signals are non-critical to function and should not be connected.

Arbitrary connection of these signals to VCC, VDDQ, VSS, or to any other signal (including each other) may result in component malfunction or incompatibility with future processors. Refer [Section 11-14, “GND, RSVD, and NCTF Signals”](#).

For reliable operation, always connect unused inputs or bi-directional signals to an appropriate signal level. Unused active high inputs should be connected through a resistor to ground (V_{SS}). Unused outputs may be left unconnected however, this may interfere with some Test Access Port (TAP) functions, complicate debug probing and prevent boundary scan testing. A resistor should be used when tying bi-directional



signals to power or ground. When tying any signal to power or ground the resistor can also be used for system testability. Resistor values should be within $\pm 20\%$ of the impedance of the baseboard trace, unless otherwise noted in the appropriate design guidelines.

Table 11-14. GND, RSVD, and NCTF Signals

Signal Name	Description
Vss	Ground: Processor ground node.
Vss_NCTF	Non-Critical To Function: These signals are for package mechanical reliability and should not be connected on the board.
RSVD	Reserved: All signals that are RSVD should not be connected on the board.
RSVD_NCTF	Reserved Non-Critical To Function: RSVD_NCTF should not be connected on the board.
RSVD_TP	Test Point: Intel recommends to route each RSVD_TP to an accessible test point. Intel may require these test points for platform specific debug. Leaving these test points inaccessible could delay debug by Intel.

11.11 Processor Internal Pull-Up / Pull-Down Terminations

Table 11-15. Processor Internal Pull-Up / Pull-Down Terminations

Signal Name	Pull Up/Pull Down	Rail	Value
BPM_N[3:0]	Pull Up/Pull Down	VCC _{IO}	16-60 Ω
PROC_PREQ#	Pull Up	VCC _{STG}	3K Ω
PROC_TDI	Pull Up	VCC _{STG}	3K Ω
PROC_TMS	Pull Up	VCC _{STG}	3K Ω
PROC_TRST#	Pull Down	VCC _{STG}	3K Ω
PROC_TCK	Pull Down	VCC _{STG}	3K Ω
CFG[19:0]	Pull Up	VCC _{IO}	3K Ω

§ §

12 Electrical Specifications

12.1 Processor Power Rails

Power Rail	Description	Y processor Line	U processor Line
V _{CCIN}	Input FIVR ¹ , Processor IA Cores And Graphic Power Rail	SVID	SVID
V _{CCIN_AUX} ⁴	Input FIVR ¹ , SA And PCH components	PCH VID	PCH VID
V _{CCST} ⁵	Sustain Power Rail	Fixed	Fixed
V _{CCSTG} ⁵	Sustain Gated Power Rail	Fixed	Fixed
V _{CCPLL}	Processor PLLs power Rail	Fixed	Fixed
V _{CCPLL_OC} ³	Processor PLLs OC power Rail	Fixed	Fixed
V _{DDQ}	Integrated Memory Controller Power Rail	Fixed (Memory technology dependent)	Fixed (Memory technology dependent)
V _{CC1P8A}	Package Rail, Type C,PCH	Fixed	Fixed
Notes: <ol style="list-style-type: none"> 1. FIVR = Fully Integrated Voltage Regulator refer Section 12.1.2, "Integrated Voltage Regulator". 2. For details regarding each rail's VR, refer to the appropriate Design Guidelines. 3. V_{CCPLL_OC} power rail should be sourced from the V_{DDQ} VR. The connection should be through a load switch in Y Processor, in U Processor the connection can be direct or through load switch depending on desired power optimization. 4. V_{CCIN_AUX} is having few point of voltage define by PCH VID. 5. V_{CCST} and V_{CCSTG} these rails are not connect to external voltage regulator moreover they are connected to the V_{CC1P05} power rail (from PCH) through a power gate. 			

12.1.1 Power and Ground Pins

All power pins should be connected to their respective processor power planes, while all VSS pins should be connected to the system ground plane. Use of multiple power and ground planes is recommended to reduce I*R drop.

12.1.2 Integrated Voltage Regulator

Due to the integration of platform voltage regulators into the processor, the processor has one main voltage rail (V_{CCIN}), the PCH has one main voltage rail (V_{CCIN_AUX}) and a voltage rail for the memory interface (V_{DDQ}).

The voltage rail V_{CCIN} will supply the integrated voltage regulators which in turn will regulate to the appropriate voltages for the Cores, cache, System Agent, TCSS and graphics. This integration allows the processor to better control on-die voltages to optimize between performance and power savings. The V_{CCIN} rail will remain a VID-based voltage with a loadline similar to the core voltage rail in previous processors.



12.1.3 V_{CC} Voltage Identification (VID)

Intel processors/chipsets are individually calibrated in the factory to operate on a specific voltage/frequency and operating-condition curve specified for that individual processor. In normal operation, the processor autonomously issues voltage control requests according to this calibrated curve using the serial voltage-identifier (SVID) interface. Altering the voltage applied at the processor/chipset causing operation outside of this calibrated curve is considered out-of-specification operation.

The SVID bus consists of three open-drain signals: clock, data, and alert# to both set voltage-levels and gather telemetry data from the voltage regulators. Voltages are controlled per an 8-bit integer value, called a VID, that maps to an analog voltage level. An offset field also exists that allows altering the VID table. Alert can be used to inform the processor that a voltage-change request has been completed or to interrupt the processor with a fault notification.

12.2 DC Specifications

The processor DC specifications in this section are defined at the processor signal pins, unless noted otherwise.

- The DC specifications for the LPDDR4/LPDDR4x/DDR4 signals are listed in the *Voltage and Current Specifications* section.
- The *Voltage and Current Specifications* section lists the DC specifications for the processor and are valid only while meeting specifications for junction temperature, clock frequency, and input voltages. Read all notes associated with each parameter.
- AC tolerances for all rails include voltage transients and voltage regulator voltage ripple up to 1 MHz. Refer additional guidance for each rail.



12.2.1 Processor Power Rails DC Specifications

12.2.1.1 V_{CCIN} DC Specifications

Table 12-1. Processor V_{CCIN} Active and Idle Mode DC Voltage and Current Specifications (Sheet 1 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
Operating Voltage	Voltage Range for Processor Operating Mode	All	0	—	2.0	V	1,2,3,7,12
I _{CCMAX} (U Processor)	Maximum Processor I _{CC}	U-Processor Line (15W) 4-Core GT2	—	—	70	A	4,6,7,11
I _{CCMAX} (U Processor)	Maximum Processor I _{CC}	U-Processor Line (15W) 2-Core GT2	—	—	55	A	4,6,7,11
I _{CCMAX} (Y Processor)	Maximum Processor I _{CC}	Y-Processor Line (9W) 4-Core GT2	—	—	49	A	4,6,7,11
I _{CCMAX} (Y Processor)	Maximum Processor I _{CC}	Y-Processor Line (9W) 2-Core GT2	—	—	35	A	4,6,7,11
I _{CC} TDC	Thermal Design Current (TDC) for processor V _{CCIN} Rail	—	—	—	—	A	9
TOB _{VCC}	Voltage Tolerance	PS0, PS1	—	—	±20	mV	3, 6, 8
		PS2, PS3	—	—	±35		
Ripple	Ripple Tolerance	PS0, PS1	—	—	±15	mV	3, 6, 8
		PS2, PS3	—	—	±30		
DC_LL	Loadline slope within the VR regulation loop capability (<=3 KHz)	U-Processor Line	0	—	2	mΩ	10,13,14,15
		Y-Processor Line	0	—	2	mΩ	10,13,14,15
AC_LL3	AC Loadline 3 (>=3 KHz)	U-Processor Line	—	—	4.2	mΩ	10,13,14
		Y-Processor Line	—	—	4.7	mΩ	10,13,14
T_OVS_TDP_MAX	Max Overshoot time TDP/virus mode	—	—	—	500	μs	
V_OVS_TDP_MAX/virus_MAX	Max Overshoot at TDP/virus mode	—	—	—	10	%	


Table 12-1. Processor V_{CCIN} Active and Idle Mode DC Voltage and Current Specifications (Sheet 2 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
Notes:							
1. Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.							
2. Each processor is programmed with a maximum valid voltage identification value (VID) that is set at manufacturing and cannot be altered. Individual maximum VID values are calibrated during manufacturing such that two processors at the same frequency may have different settings within the VID range. Note that this differs from the VID employed by the processor during a power management event (Adaptive Thermal Monitor, Enhanced Intel SpeedStep Technology, or low-power states).							
3. The voltage specification requirements are measured across V _{CC_SENSE} and V _{SS_SENSE} as near as possible to the processor. measurement needs to be performed with a 20MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1Mohm minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.							
4. Processor V _{CCIN} VR to be designed to electrically support this current.							
5. Processor V _{CCIN} VR to be designed to thermally support this current indefinitely.							
6. Long term reliability cannot be assured if tolerance, ripple, and core noise parameters are violated.							
7. Long term reliability cannot be assured in conditions above or below Max/Min functional limits.							
8. PSx refers to the voltage regulator power state as set by the SVID protocol.							
9. Refer Intel Platform Design Studio (iPDS) for the minimum, typical, and maximum V _{CC} allowed for a given current and Thermal Design Current (TDC).							
10. LL measured at sense points.							
11. Typ column represents I _{CCMAX} for commercial application it is NOT a specification - it's a characterization of limited samples using limited set of benchmarks that can be exceeded.							
12. Operating voltage range in steady state.							
13. LL spec values should not be exceeded. If exceeded, power, performance and reliability penalty are expected.							
14. Load Line (AC/DC) should be measured by the VRTT tool and programmed accordingly via the BIOS Load Line override setup options. AC/DC Load Line BIOS programming directly affects operating voltages (AC) and power measurements (DC). A superior board design with a shallower AC Load Line can improve on power, performance and thermals compared to boards designed for POR impedance.							
15. Optimal value will depend on platform VR design and workload.							

12.2.1.2 V_{CC1p8A} DC Specifications

Table 12-2. Processor V_{CC1p8A} Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Notes ^{1,2}
V _{CC1p8A}	Package voltage (DC specification)	All	—	1.8	—	V	1,3
I _{CCMAX_1p8A}	Max Current for 1p8A Rail	U-Processor Line	—	—	700	mA	1
		Y-Processor Line	—	—	500	mA	
TOB V _{CC1p8A}	V _{CC1p8A} Tolerance	All	AC+DC: ± 5%			%	1,3,4
Ripple	Ripple Tolerance	All	—	—	90	mV	1
Notes:							
1. Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.							
2. Long term reliability cannot be assured in conditions above or below Max/Min functional limits.							
3. The voltage specification requirements are measured across V _{CC_SENSE} and V _{SS_SENSE} as near as possible to the processor. measurement needs to be performed with a 20MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1Mohm minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.							
4. For Voltage less than 1 VTOB will be 50 mv.							

12.2.1.3 V_{CCIN_AUX} DC Specifications

Table 12-3. V_{CCIN_AUX} Supply DC Voltage and Current Specifications (Sheet 1 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
V _{CCin_AUX}		U-Processor Line	0	1.8	—	V	1,3,4
		Y-Processor Line	0	1.65	1.8	V	1,3,4



12.2.1.3 VccIN_AUX DC Specifications

Table 12-3. VccIN_AUX Supply DC Voltage and Current Specifications (Sheet 2 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹	
Icc _{MAX}	Maximum VccIN_AUX Icc	U-Processor Line (15W) 4-Core GT2	0	—	32	A	1	
		U-Processor Line (15W) 2-Core GT2	0	—	32			
		Y-Processor Line (9W) 4-Core GT2	0	—	22			
TOB _{VCC}	Voltage Tolerance Budget	U -Processor Line	—	—	AC+DC: -10/+5		%	1,3,6
		Y-Processor Line	—	—	AC+DC: ± 7.5		%	1,3,6
VOS	Overshoot Voltage	All	—	—	—	1.95	V	7
TVOS	Overshoot Time	All	—	—	—	5	us	7
AC_LL	AC Loadline 3 (<1 MHz)	Y-Processor Line	—	—	5.9		mΩ	4,5
		U-Processor Line	—	—	4.9			
	AC Loadline 2 (1-40 MHz)	Y-Processor Line	—	—	6.5			
		U-Processor Line	—	—	8.0			
Notes:								
1. All specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.								
2. Long term reliability cannot be assured in conditions above or below Max/Min functional limits.								
3. The voltage specification requirements are measured across Vcc_SENSE and Vss_SENSE as near as possible to the processor. measurement needs to be performed with a 20 MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1Mohm minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.								
4. Max impedance allowed between 1 MHz-40 MHz is lower than LL3. Comply with recommended impedance target to avoid coupling noise concerns.								
5. The LL3 values are for reference. must still meet voltage tolerance specification.								
6. Voltage Tolerance budget values Includes ripples.								
7. Overshoot with max voltage of 2.13 V is allowed if it sustained for less then 500us.								
8. This rail can be connected to 1.65 V.								
9. VCCIN_AUX is having few point of voltage define by PCH VID.								

12.2.1.4 VDDQ DC Specifications

Table 12-4. Memory Controller (VDDQ) Supply DC Voltage and Current Specifications (Sheet 1 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
V _{DDQ} (LPDDR4/x)	Processor I/O supply voltage for LPDDR4/x	All	Typ-5%	1.1	Typ+5%	V	3,4,5
V _{DDQ} (DDR4)	Processor I/O supply voltage for DDR4	All	Typ-5%	1.2	Typ+5%	V	3,4,5
TOB _{VDDQ}	VDDQ Tolerance	All	AC+DC: ± 5%			%	3,4,6
Icc _{MAX_VDDQ} (LPDDR4/x)	Max Current for V _{DDQ} Rail (LPDDR4/x)	Y-Processor Line	—	—	3	A	2
		U-Processor Line	—	—	3.5		
Icc _{MAX_VDDQ} (DDR4)	Max Current for V _{DDQ} Rail (DDR4)	U-Processor Line	—	—	3.5		


Table 12-4. Memory Controller (V_{DDQ}) Supply DC Voltage and Current Specifications (Sheet 2 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
Notes:							
1. Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.							
2. The current supplied to the DIMM modules is not included in this specification.							
3. Includes AC and DC error, where the AC noise is bandwidth limited to under 100 MHz, measured on package pins.							
4. No requirement on the breakdown of AC versus DC noise.							
5. The voltage specification requirements are measured across V_{CC_SENSE} and V_{SS_SENSE} as near as possible to the processor. measurement needs to be performed with a 20 MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 Mohm minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.							
6. For Voltage less than 1 VTOB will be 50 mv.							

12.2.1.5 V_{CCST} DC Specifications

Table 12-5. V_{CCST} Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Units	Notes ^{1,2}
V_{CCST}	Processor Vcc Sustain supply voltage	All Processor Lines	—	1.025	—	V	3
TOB_{ST}	V_{CCST} Tolerance	All	AC+DC:± 5%			%	3,5
ICC_{MAX_ST}	Max Current for V_{CCST}	U-Processor Line	—	—	800	mA	4
		Y-Processor Line	—	—	300		
Notes:							
1. Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.							
2. Long term reliability cannot be assured in conditions above or below Max/Min functional limits.							
3. The voltage specification requirements are measured across V_{CC_SENSE} and V_{SS_SENSE} as near as possible to the processor. measurement needs to be performed with a 20 MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 Mohm minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.							
4. The maximum ICC_{MAX_ST} specification is preliminary and based on initial pre-silicon estimation and is subject to change.							
5. For Voltage less than 1 V TOB will be 50 mv.							

Table 12-6. V_{CCSTG} Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Units	Notes ^{1,2}
V_{CCSTG}	Processor Vcc Sustain supply voltage	All	—	1.025	—	V	3
TOB_{STG}	V_{CCSTG} Tolerance	All	AC+DC:± 5%			%	3,5
ICC_{MAX_STG}	Max Current for V_{CCSTG}	U -Processor Line	—	—	150	mA	4
		Y-Processor Line	—	—	60		
Notes:							
1. All specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.							
2. Long term reliability cannot be assured in conditions above or below Max/Min functional limits.							
3. The voltage specification requirements are measured across V_{CC_SENSE} and V_{SS_SENSE} as near as possible to the processor. measurement needs to be performed with a 20 MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 Mohm minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.							
4. The maximum ICC_{MAX_ST} specification is preliminary and based on initial pre-silicon estimation and is subject to change.							
5. For Voltage less than 1 V TOB will be 50 mv.							



12.2.1.6 V_{CCPLL} DC Specifications

Table 12-7. Processor PLL (V_{CCPLL}) Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Notes ^{1,2}
V _{CCPLL}	PLL supply voltage (DC specification)	All	—	1.025	—	V	3
TOB _{CCPLL}	V _{CCPLL_OC} Tolerance	All	AC+DC: ± 5%			%	3,4
I _{CCMAX_VCCPLL}	Max Current for V _{CCPLL} Rail	U-Processor Line	—	—	90	mA	
		Y-Processor Line	—	—			
Notes: <ol style="list-style-type: none"> Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date. Long term reliability cannot be assured in conditions above or below Max/Min functional limits. The voltage specification requirements are measured across V_{CC_SENSE} and V_{VSS_SENSE} as near as possible to the processor. measurement needs to be performed with a 20 MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 Mohm minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe. For Voltage less than 1 V TOB will be 50mv. 							

Table 12-8. Processor PLL_OC (V_{CCPLL_OC}) Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Notes ^{1,2}
V _{CCPLL_OC}	PLL_OC supply voltage (DC specification)	All	—	V _{DDQ}	—	V	3
TOB _{CCPLL_OC}	V _{CCPLL_OC} Tolerance	All	AC+DC: ± 5%			%	3,4
I _{CCMAX_VCCPLL_OC}	Max Current for V _{CCPLL} Rail	U-Processor Line	—	—	160	mA	5
		Y-Processor Line	—	—	170		
Notes: <ol style="list-style-type: none"> Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date. Long term reliability cannot be assured in conditions above or below Max/Min functional limits. The voltage specification requirements are measured across V_{CC_SENSE} and V_{VSS_SENSE} as near as possible to the processor. measurement needs to be performed with a 20 MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 Mohm minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe. For Voltage less than 1 V TOB will be 50 mv. I_{CCMAX} values rely on voltage V_{DDQ}=1.1 V. 							



12.2.2 Processor Interfaces DC Specifications

12.2.2.1 DDR4 DC Specifications

Table 12-9. DDR4 Signal Group DC Specifications (Sheet 1 of 2)

Symbol	Parameter	U-Processor Line			Units	Notes ¹
		Minimum	Typical	Maximum		
V _{IL}	Input Low Voltage	—	0.75* V _{DDQ}	0.68* V _{DDQ}	V	2, 3, 4
V _{IH}	Input High Voltage	0.82* V _{DDQ}	0.75* V _{DDQ}	—	V	2, 3, 4
R _{ON_UP} (DQ)	Data Buffer pull-up Resistance	25	—	60	Ω	5, 12
R _{ON_DN} (DQ)	Data Buffer pull-down Resistance	26	—	75		
R _{ODT} (DQ)	On-die termination equivalent resistance for data signals	25	—	Hi-Z	Ω	6, 12
V _{ODT} (DC)	On-die termination DC working point (driver set to receive mode)	0.7* V _{DDQ}	0.75* V _{DDQ}	0.8* V _{DDQ}	V	12
R _{ON_UP} (CK)	Clock Buffer pull-up Resistance	25	—	60	Ω	5, 12
R _{ON_DN} (CK)	Clock Buffer pull-down Resistance	25	—	75	Ω	5, 12
R _{ON_UP} (CMD)	Command Buffer pull-up Resistance	23	—	50	Ω	5, 12
R _{ON_DN} (CMD)	Command Buffer pull-down Resistance	24	—	57	Ω	5, 12
R _{ON_UP} (CTL)	Control Buffer pull-up Resistance	23	—	50	Ω	5, 12
R _{ON_DN} (CTL)	Control Buffer pull-down Resistance	24	—	57	Ω	5, 12
R _{ON_UP} (SM_PG_CNTL1)	System Memory Power Gate Control Buffer Pull-up Resistance	45	—	125	Ω	—
R _{ON_DN} (SM_PG_CNTL1)	System Memory Power Gate Control Buffer Pull-down Resistance	40	—	130	Ω	—
I _{LI}	Input Leakage Current (DQ, CK) 0 V 0.2*V _{DDQ} 0.8*V _{DDQ}	—	—	1.1	mA	—
DDR0_VREF_DQ DDR1_VREF_DQ DDR_VREF_CA	VREF output voltage	Trainable	V _{DDQ} /2	Trainable	V	—
SM_RCOMP[0]	Command COMP Resistance	99	100	101	Ω	8
SM_RCOMP[1]	Data COMP Resistance	99	100	101	Ω	8
SM_RCOMP[2]	ODT COMP Resistance	99	100	101	Ω	8



Table 12-9. DDR4 Signal Group DC Specifications (Sheet 2 of 2)

Symbol	Parameter	U-Processor Line			Units	Notes ¹
		Minimum	Typical	Maximum		
Notes:						
1. All specifications in this table apply to all processor frequencies. Timing specifications only depend on the operating frequency of the memory channel and not the maximum rated frequency. 2. V_{IL} is defined as the maximum voltage level at a receiving agent that will be interpreted as a logical low value. 3. V_{IH} is defined as the minimum voltage level at a receiving agent that will be interpreted as a logical high value. 4. V_{IH} and V_{OH} may experience excursions above V_{DDQ} . However, input signal drivers should comply with the signal quality specifications. 5. Pull up/down resistance after compensation (assuming $\pm 5\%$ COMP inaccuracy). 6. BIOS power training may change these values significantly based on margin/power trade-off. 7. ODT values after COMP (assuming $\pm 5\%$ inaccuracy). BIOS MRC can reduce ODT strength towards. 8. The minimum and maximum values for these signals are programmable by BIOS to one of the two sets. 9. SM_RCOMP[x] resistance should be provided on the system board with 1% resistors. SM_RCOMP[x] resistors are to VSS. Values are pre-silicon estimations and are subject to change. 10. SM_DRAMPWROK must have a maximum of 15 ns rise or fall time over $V_{DDQ} * 0.30 \pm 100$ mV and the edge must be monotonic. 11. SM_VREF is defined as $V_{DDQ}/2$ for DDR4/LPDDR4. 12. R_{ON} tolerance is preliminary and might be subject to change. 13. Max-min range is correct but center point is subject to change during MRC boot training. 14. Processor may be damaged if V_{IH} exceeds the maximum voltage for extended periods.						

12.2.2.2 LPDDR4/x DC Specifications

Table 12-10. LPDDR4/x Signal Group DC Specifications (Sheet 1 of 2)

Symbol	Parameter	U/Y-Processor Line			Units	Notes ¹
		Minimum	Typical	Maximum		
V_{IL}	Input Low Voltage	—	0.2* V_{DDQ}	0.08* V_{DDQ}	V	2, 3, 4
V_{IH}	Input High Voltage	0.35* V_{DDQ}	0.2* V_{DDQ}	—	V	2, 3, 4
$R_{ON_UP(DQ)}$	Data Buffer pull-up Resistance	25 (LP4x:23)	—	60 (LP4x:58)	Ω	5, 12
$R_{ON_DN(DQ)}$	Data Buffer pull-down Resistance	25 (LP4x:26)	—	72 (LP4x:85)	Ω	5, 12
$R_{ODT(DQ)}$	On-die termination equivalent resistance for data signals	28 (LP4x:26)	—	Hi-Z	Ω	6, 12
$V_{ODT(DC)}$	On-die termination DC working point (driver set to receive mode)	0.15* V_{DDQ} (LP4x: 0.25* V_{DDQ})	0.2* V_{DDQ} (LP4x: 0.3* V_{DDQ})	0.25* V_{DDQ} (LP4x:0.35* V_{DDQ})	V	10
$R_{ON_UP(CK)}$	Clock Buffer pull-up Resistance	24 (LP4x:30)	—	60 (LP4x:59)	Ω	5, 12
$R_{ON_DN(CK)}$	Clock Buffer pull-down Resistance	28	—	92 (LP4x:94)	Ω	5, 12
$R_{ON_UP(CMD)}$	Command Buffer pull-up Resistance	26	—	50	Ω	5, 12
$R_{ON_DN(CMD)}$	Command Buffer pull-down Resistance	22 (LP4x:20)	—	67	Ω	5, 12
$R_{ON_UP(CTL)}$	Control Buffer pull-up Resistance	26	—	50	Ω	5, 12
$R_{ON_DN(CTL)}$	Control Buffer pull-down Resistance	22 (LP4x:20)	—	67	Ω	5, 12



Table 12-10. LPDDR4/x Signal Group DC Specifications (Sheet 2 of 2)

Symbol	Parameter	U/Y-Processor Line			Units	Notes ¹
		Minimum	Typical	Maximum		
R _{ON_UP} (SM_VTT_CTL1)	System Memory Power Gate Control Buffer Pull-up Resistance	N/A	—	N/A	Ω	N/A
R _{ON_DN} (SM_VTT_CTL1)	System Memory Power Gate Control Buffer Pull-down Resistance	N/A	—	N/A	Ω	N/A
I _{LI}	Input Leakage Current (DQ, CK) 0 V 0.2*V _{DDQ} 0.8*V _{DDQ}	—	—	1	mA	—
DDR0_VREF_DQ DDR1_VREF_DQ DDR_VREF_CA	VREF output voltage	Trainable			V	—
SM_RCOMP[0]	Command COMP Resistance	99	100	101	Ω	8
SM_RCOMP[1]	Data COMP Resistance	99	100	101	Ω	8
SM_RCOMP[2]	ODT COMP Resistance	99	100	101	Ω	8
Notes:						
1. All specifications in this table apply to all processor frequencies. Timing specifications only depend on the operating frequency of the memory channel and not the maximum rated frequency.						
2. V _{IL} is defined as the maximum voltage level at a receiving agent that will be interpreted as a logical low value.						
3. V _{IH} is defined as the minimum voltage level at a receiving agent that will be interpreted as a logical high value.						
4. V _{IH} and V _{OH} may experience excursions above V _{DDQ} . However, input signal drivers should comply with the signal quality specifications.						
5. Pull up/down resistance after compensation (assuming ±5% COMP inaccuracy). Note that BIOS power training may change these values significantly based on margin/power trade-off.						
6. ODT values after COMP (assuming ±5% inaccuracy). BIOS MRC can reduce ODT strength towards						
7. The minimum and maximum values for these signals are programmable by BIOS to one of the two sets.						
8. SM_RCOMP[x] resistance should be provided on the system board with 1% resistors. SM_RCOMP[x] resistors are to VSS. Values are pre-silicon estimations and are subject to change.						
9. SM_DRAMPWROK must have a maximum of 15 ns rise or fall time over VDDQ * 0.30 ±100 mV and the edge must be monotonic.						
10. SM_VREF is defined as VDDQ/2 for DDR4/LPDDR4.						
11. R _{ON} tolerance is preliminary and might be subject to change.						
12. Max-min range is correct but center point is subject to change during MRC boot training.						
13. Processor may be damaged if V _{IH} exceeds the maximum voltage for extended periods.						

12.2.2.3 Digital Display Interface (DDI) DC Specifications

Table 12-13. Digital Display Interface Group DC Specifications (DP/HDMI) (Sheet 1 of 2)

Symbol	Parameter	Minimum	Typical	Maximum	Units	Notes ¹
V _{IL}	Aux Input Low Voltage	—	—	0.8	V	
V _{IH}	Aux Input High Voltage	2.25	—	3.6	V	
V _{OL}	DDIB_TXC[3:0] Output Low Voltage DDIC_TXC[3:0] Output Low Voltage DDID_TXC[3:0] Output Low Voltage	—	—	0.25*V _{CCIO}	V	1,2
V _{OH}	DDIB_TXC[3:0] Output High Voltage DDIC_TXC[3:0] Output High Voltage DDID_TXC[3:0] Output High Voltage	0.75*V _{CCIO}	—	—	V	1,2
Z _{TX-DIFF-DC}	DC Differential Tx Impedance	100	—	120	Ω	



Table 12-13. Digital Display Interface Group DC Specifications (DP/HDMI) (Sheet 2 of 2)

Symbol	Parameter	Minimum	Typical	Maximum	Units	Notes ¹
Notes:						
1. V_{CCIO} depends on segment.						
2. V_{OL} and V_{OH} levels depends on the level chosen by the Platform.						

12.2.2.4 embedded DisplayPort* (eDP*) DC Specification

Table 12-14. embedded DisplayPort* (eDP*) Group DC Specifications

Symbol	Parameter	Minimum	Typical	Maximum	Units
V_{OL}	eDP_DISP_UTIL Output Low Voltage	—	—	$0.1*V_{CCIO}$	V
V_{OH}	eDP_DISP_UTIL Output High Voltage	$0.9*V_{CCIO}$	—	—	V
R_{UP}	eDP_DISP_UTIL Internal pull-up	45	—	—	Ω
R_{DOWN}	eDP_DISP_UTIL Internal pull-down	45	—	—	Ω

12.2.2.5 MIPI* CSI-2 D-Phy Receiver DC Specifications

Symbol	Parameter	Minimum	Typical	Maximum	Units	Notes
$V_{CMRX(DC)}$	Common-mode voltage HS receive mode	70	—	330	mV	1,2
V_{IDTH}	Differential input high threshold	—	—	70	mV	3
		—	—	40	mV	4
V_{IDTL}	Differential input low threshold	-70	—	—	mV	3
		-40	—	—	mV	4
V_{IHHS}	Single-ended input high voltage	—	—	460	mV	1
V_{ILHS}	Single-ended input low voltage	-40	—	—	mV	1
$V_{TERM-EN}$	Single-ended threshold for HS termination enable	—	—	450	mV	
Z_{ID}	Differential input impedance	80	100	125	Ω	
Notes:						
1. Excluding possible additional RF interference of 100 mV peak sine wave beyond 450 MHz.						
2. This table value includes a ground difference of 50 mV between the transmitter and the receiver, the static common-mode level tolerance and variations below 450 MHz.						
3. For devices supporting data rates < 1.5 Gbps.						
4. For devices supporting data rates > 1.5 Gbps.						
5. Associated Signals: MIPI* CSI2: Refer to MIPI® Alliance D-PHY Specification 1.2.						

12.2.2.6 CMOS DC Specifications

Table 12-15. CMOS Signal Group DC Specifications (Sheet 1 of 2)

Symbol	Parameter	Minimum	Maximum	Units	Notes ¹
V_{IL}	Input Low Voltage	—	$V_{CC}*0.3$	V	2, 5
V_{IH}	Input High Voltage	$V_{CC}*0.7$	—	V	2, 4, 5
R_{ON}	Buffer on Resistance	20	70	Ω	-
I_{LI}	Input Leakage Current	—	± 150	μA	3



Table 12-15. CMOS Signal Group DC Specifications (Sheet 2 of 2)

Symbol	Parameter	Minimum	Maximum	Units	Notes ¹
Notes: 1. Unless otherwise noted, all specifications in this table apply to all processor frequencies. 2. The Vcc referred to in these specifications refers to instantaneous Vcc _{ST/IO} . 3. For V _{IN} between "0" V and Vcc _{ST} . Measured when the driver is tri-stated. 4. V _{IH} may experience excursions above Vcc _{ST} . However, input signal drivers should comply with the signal quality specifications. 5. N/A.					

12.2.2.7 GTL and OD DC Specification

Table 12-16. GTL Signal Group and Open Drain Signal Group DC Specifications

Symbol	Parameter	Minimum	Maximum	Units	Notes ¹
V _{IL}	Input Low Voltage (TAP, except PROC_JTAG_TCK, PROC_JTAG_TRST#)	—	0.6*Vcc	V	2, 5
V _{IH}	Input High Voltage (TAP, except PROC_JTAG_TCK, PROC_JTAG_TRST#)	0.72*Vcc	—	V	2, 4, 5
V _{IL}	Input Low Voltage (PROC_JTAG_TCK, PROC_JTAG_TRST#)	—	0.3*Vcc	V	2, 5
V _{IH}	Input High Voltage (PROC_JTAG_TCK, PROC_JTAG_TRST#)	0.7*Vcc	—	V	2, 4, 5
V _{HYSTERESIS}	Hysteresis Voltage	0.2*Vcc	—	V	-
R _{ON}	Buffer on Resistance (TDO)	7	17	Ω	-
V _{IL}	Input Low Voltage (other GTL)	—	0.6*Vcc	V	2, 5
V _{IH}	Input High Voltage (other GTL)	0.72*Vcc	—	V	2, 4, 5
R _{ON}	Buffer on Resistance (BPM)	12	28	Ω	-
R _{ON}	Buffer on Resistance (other GTL)	16	24	Ω	-
I _{LI}	Input Leakage Current	—	±250	μA	3
Notes: 1. Unless otherwise noted, all specifications in this table apply to all processor frequencies. 2. The Vcc referred to in these specifications refers to instantaneous Vcc _{ST/IO} . 3. For V _{IN} between 0V and Vcc. Measured when the driver is tri-stated. 4. V _{IH} and V _{OH} may experience excursions above Vcc. However, input signal drivers should comply with the signal quality specifications. 5. N/A.					

12.2.2.8 PECCI DC Characteristics

The PECCI interface operates at a nominal voltage set by Vcc_{ST}. The set of DC electrical specifications shown in the following table is used with devices normally operating from a Vcc_{ST} interface supply.

Vcc_{ST} nominal levels will vary between processor families. All PECCI devices will operate at the Vcc_{ST} level determined by the processor installed in the system.

Table 12-17. PECCI DC Electrical Limits (Sheet 1 of 2)

Symbol	Definition and Conditions	Minimum	Maximum	Units	Notes ¹
R _{up}	Internal pull up resistance	15	45	Ω	3
V _{in}	Input Voltage Range	-0.15	Vcc _{ST} + 0.15	V	-
V _{hysteresis}	Hysteresis	0.1 * Vcc _{ST}	—	V	-

Table 12-17. PECE DC Electrical Limits (Sheet 2 of 2)

Symbol	Definition and Conditions	Minimum	Maximum	Units	Notes ¹
V _{IL}	Input Voltage Low- Edge Threshold Voltage	0.275 * V _{CCST}	0.525 * V _{CCST}	V	-
V _{IH}	Input Voltage High- Edge Threshold Voltage	0.550 * V _{CCST}	0.725 * V _{CCST}	V	-
C _{bus}	Bus Capacitance per Node	—	10	pF	-
C _{pad}	Pad Capacitance	0.7	1.8	pF	-
I _{leak000}	leakage current @ 0V	—	0.25	mA	-
I _{leak100}	leakage current @ V _{CCST}	—	0.15	mA	-


Notes:

- V_{CCST} supplies the PECE interface. PECE behavior does not affect V_{CCST} min/max specifications.
- The leakage specification applies to powered devices on the PECE bus.
- The PECE buffer internal pull up resistance measured at 0.75* V_{CCST}.

Input Device Hysteresis

The input buffers in both client and host models should use a Schmitt-triggered input design for improved noise immunity. Use the following figure as a guide for input buffer design.

Figure 12-1. Input Device Hysteresis



12.3 Test Access Port (TAP) Connection

Due to the voltage levels supported by other components in the Test Access Port (TAP) logic, Intel recommends the processor be first in the TAP chain, followed by any other components within the system. A translation buffer should be used to connect to the rest of the chain unless one of the other components is capable of accepting an input of the appropriate voltage. Two copies of each signal may be required with each driving a different voltage level.

The processor supports Boundary Scan (JTAG) IEEE 1149.1-2001 and IEEE 1149.6-2003 standards.





13 Package Mechanical Specifications

13.1 Package Mechanical Attributes

The U/Y Processor Line use a Flip Chip technology available in a Ball Grid Array (BGA) package. The following table provides an overview of the of the mechanical attributes of the package.

Table 13-1. Package Mechanical Attributes

Package	Parameter	Y-Processor Line	U- Processor Line
		4 Core GT2	4/2 Core GT2
Package Technology	Package Type	Flip Chip Ball Grid Array	Flip Chip Ball Grid Array
	Interconnect	Ball Grid Array (BGA)	Ball Grid Array (BGA)
	Lead Free	Yes	Yes
	Halogenated Flame Retardant Free	Yes	Yes
Package Configuration	Solder Ball Composition	SAC405	SAC405
	Ball/Pin Count	1377	1526
	NCTF Corner balls	11 balls per corner, 9@A1	4 to 6 balls per corner
	Grid Array Pattern	Balls Anywhere	Balls Anywhere
	Land Side Capacitors	Yes (250 um max height)	Yes (250 um max height)
	Die Side Capacitors	No	No
	Die Configuration	2 Dice Multi-Chip Package (MCP)	2 Dice Multi-Chip Package (MCP)
Package Dimensions	Nominal Package Size	26.5x18.5 mm	50x25 mm
	Z-height	0.936 ± 0.088	1.197 ± 0.096
	Min Ball/Pin pitch	0.43 mm	0.65 mm

13.2 Package Loading and Die Pressure Specifications

Intel has defined the maximum total compressive load limits that can be applied to the package for the following SKUs. This value should not be exceeded by the system design.



13.2.1 Package Loading Specifications

Package	Maximum Static Normal Load (preliminary data)	Backing Plate Assumptions	Minimum PCB Thickness Assumptions [mm/mils]	Notes
Y-Processor Line	10	NO	0.7-0.9\28-36	1,2,3,6,7,8,9
	5	NO	0.6 \ 24	1,2,3,6,7,8,9
U-Processor Line	15	NO	0.8-1.2 \32-47	1,2,3,5,6,7,8,9
Notes: <ol style="list-style-type: none"> The thermal solution attach mechanism should not induce continuous stress to the package. It may only apply a uniform load to the die to maintain a thermal interface. This specification applies to the uniform compressive load in the direction perpendicular to the dies' top surface. Load should be centered on processor die center. This specification is based on limited testing for design characterization. All values are pre-silicon values and are subject to change. Backing plate is also acceptable if desired. Considerations should be made to ensure steady state static loading on the packages does not exceed the limits recommended. Excessive Steady State static loading can induce solder ball cracks especially over a period of time resulting in higher failure rate. This static compressive load is not to be exceeded, therefore the tolerance of the package and the tolerances of the thermal solution (including attach mechanism) should be taken into account when calculating or measuring static load on the package. An ideal thermal solution design would apply a load as uniform as possible on all dies in order to optimize thermal performance and minimize mechanical risk. Thermal structural support should be attached to the motherboard (as a backing plate or block) or built into the system base, when applicable. 				

13.2.2 Die Pressure Specifications

A more relevant metric for concentrated loading is chosen by Intel based on the physics of failure to evaluate die damage risk due to thermal solution enabling.

Static Compressive pressure refers to the long term steady state pressure applied to the die from the thermal solution after system assembly is complete.

Transient Compressive pressure refers to the pressure on the dice at any moment during the thermal solution assembly/disassembly procedures. Other system procedures such as repair/rework can also cause high pressure loading to occur on the die and should be evaluated to ensure these limits are not exceeded.

Metric: This metric is pressure over a 2 mm x 2 mm area.

Table 13-2. Package Loading Specifications

Package	Static Compressive Pressure ¹ [PSI]	Transient Compressive Pressure ¹ [PSI]
Y-Processor Line	800	800
U-Processor Line	800	800
Note: This is the load and pressure that has been tested by Intel for a single assembly cycle. This metric is pressure over 2 mm ² (2 mm x 2 mm) area.		



13.3 Package Storage Specifications

Parameter	Description	Min	Max	Notes
T _{ABSOLUTE STORAGE}	The non-operating device storage temperature. Damage (latent or otherwise) may occur when subjected to this temperature for any length of time in Intel Original sealed moisture barrier bag and / or box.	-25 °C	125 °C	1, 2, 3
T _{SUSTAINED STORAGE}	The ambient storage temperature limit (in shipping media) for the sustained period of time as specified below in Intel Original sealed moisture barrier bag and / or box.	-5 °C	40 °C	1, 2, 3
RH _{SUSTAINED STORAGE}	The maximum device storage relative humidity for the sustained period of time as specified below in Intel Original sealed moisture barrier bag and / or box.	60% @ 24 °C		1, 2, 3
TIME _{SUSTAINED STORAGE}	Maximum time: associated with customer shelf life in Intel Original sealed moisture barrier bag and / or box.	NA	Moisture Sensitive Devices: 60 months from bag seal date; Non-moisture sensitive devices: 60 months from lot date	1, 2, 3
Storage Conditions	Processors in a non-operational state may be installed in a platform, in a tray, boxed, or loose and may be sealed in airtight package or exposed to free air. Under these conditions, processor landings should not be connected to any supply voltages, have any I/Os biased, or receive any clocks. Upon exposure to "free air" (that is, unsealed packaging or a device removed from packaging material), the processor should be handled in accordance with moisture sensitivity labeling (MSL) as indicated on the packaging material. Boxed Land Grid Array packaged (LGA) processors are MSL 1 ('unlimited' or unaffected) as they are not heated in order to be inserted in the socket.			
Notes:				
<ol style="list-style-type: none"> 1. T_{ABSOLUTE STORAGE} applies to the un-assembled component only and does not apply to the shipping media, moisture barrier bags or desiccant. Refers to a component device that is not assembled in a board or socket that is not to be electrically connected to a voltage reference or I/O signals. 2. Specified temperatures are based on data collected. The JEDEC, J-STD-020 moisture level rating and associated handling practices apply to all moisture sensitive devices removed from the moisture barrier bag. 3. Post board attach storage temperature limits are not specified for non-Intel branded boards. Consult the board manufacturer for storage specifications. 				





14 CPU And Device IDs

14.1 CPUID

The processor ID and stepping can be identified by the following register contents:

Table 14-1. CPUID Format

Field	Reserved	Extended Family	Extended Model	Reserved	Processor Type	Family Code	Model Number	Stepping ID
Bits	31:28	27:20	19:16	15:14	13:12	11:8	7:4	3:0

Table 14-2. Component Identification

SKU	CPUID
Y/U	0x706E5

- The Extended Family, Bits [27:20] are used in conjunction with the Family Code, specified in Bits[11:8], to indicate whether the processor belongs to the Intel386™, Intel486™, Pentium®, Pentium 4, or Intel® Core™ processor family.
- The Extended Model, Bits [19:16] in conjunction with the Model Number, specified in Bits [7:4], are used to identify the model of the processor within the processor's family.
- The Family Code corresponds to Bits [11:8] of the EDX register after RESET, Bits [11:8] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
- The Model Number corresponds to Bits [7:4] of the EDX register after RESET, Bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.
- The Stepping ID in Bits [3:0] indicates the revision number of that model.
- When EAX is initialized to a value of '1', the CPUID instruction returns the Extended Family, Extended Model, Processor Type, Family Code, Model Number and Stepping ID value in the EAX register. Note that the EDX processor signature value after reset is equivalent to the processor signature output value in the EAX register.

Cache and TLB descriptor parameters are provided in the EAX, EBX, ECX and EDX registers after the CPUID instruction is executed with a 2 in the EAX register.



14.2 PCI Configuration Header

Every PCI-compatible function has a standard PCI configuration header, as shown in Table 14-3, “PCI Configuration Header”. This includes mandatory registers (Bold) to determine which driver to load for the device. Some of these registers define ID values for the PCI function, which are described in this chapter.

Table 14-3. PCI Configuration Header

Byte3	Byte2	Byte1	Byte0	Address
Device ID		Vendor ID (0x8086)		00h
Status		Command		04h
Class Code			Revision ID	08h
BIST	Header Type	Latency Timer	Cache Line Size	0Ch
Base Address Register0 (BAR0)				10h
Base Address Register1 (BAR1)				14h
Base Address Register2 (BAR2)				18h
Base Address Register3 (BAR3)				1Ch
Base Address Register4 (BAR4)				20h
Base Address Register5 (BAR5)				24h
Subsystem ID (0x7270)		Subsystem Vendor ID (0x8086)		28h
Expansion ROM Base Address				2Ch
Reserved			Capabilities Pointer	30h
Reserved				34h
Max Latency	Min Grant	Interrupt Pin	Interrupt Line	3Ch

14.3 Device IDs

Table 14-4. Host Device ID (DID0)

Platform	Device ID
Y Processor 2 Cores	0x8A00h
U Processor 2 Cores	0x8A02h
Y Processor 4 Cores	0x8A10h
U Processor 4 Cores	0x8A12h

Table 14-5. Other Device ID (Sheet 1 of 2)

Device	Processor Line	Bus / Device / Function	DID
Graphics	All	0 / 2 / 0	-
IPU	Y, U	0 / 5 / 0	0x8A19
TBT_PCIe0	All	0 / 7 / 0	0x8A1D
TBT_PCIe1	All	0 / 7 / 1	0x8A1F
TBT_PCIe2	All	0 / 7 / 2	0x8A21
TBT_PCIe3	All	0 / 7 / 3	0x8A23
GNA	All	0 / 8 / 0	0x8A11



Table 14-5. Other Device ID (Sheet 2 of 2)

Device	Processor Line	Bus / Device / Function	DID
ITH	All	0 / 9 / 0	0x8A29
USB xHCI	Y, U	0 / 13 / 0	0x8A13
USB xDCI	Y, U	0 / 13 / 1	0x8A15
TBT DMA0	All	0 / 13 / x [2-7]	0x8A0D
TBT DMA1	All	0 / 13 / x [2-7]	0x8A17

§ §