


Intel[®] Xeon[®] Processor E3-1200 v4 Product Family

Datasheet – Volume 1 of 2

June 2015



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at Intel.com, or from the OEM or retailer.

No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses.

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.

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.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

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.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel, Intel Core, Intel Xeon, Intel® High Definition Audio, Intel® Advanced Vector Extensions, Enhanced Intel Speedstep® Technology, 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 © 2014-2015, Intel Corporation. All Rights Reserved.



Table of Contents


1	Introduction	11
1.1	Supported Technologies	12
1.2	Interfaces	13
1.3	Power Management Support	13
1.3.1	Processor Core	13
1.3.2	System	13
1.3.3	Memory Controller	13
1.3.4	PCI Express*	13
1.3.5	DMI	13
1.3.6	Processor Graphics Controller	14
1.4	Thermal Management Support	14
1.5	Package Support	14
1.6	Processor Testability	14
1.7	Terminology	14
1.8	Related Documents	18
2	Interfaces	19
2.1	System Memory Interface	19
2.1.1	System Memory Technology Supported	20
2.1.2	System Memory Timing Support	21
2.1.3	System Memory Organization Modes	21
2.1.3.1	System Memory Frequency	23
2.1.3.2	Intel® Fast Memory Access (Intel® FMA) Technology Enhancements	23
2.1.3.3	Data Scrambling	23
2.2	PCI Express* Interface	24
2.2.1	PCI Express* Support	24
2.2.2	PCI Express* Architecture	25
2.2.3	PCI Express* Configuration Mechanism	26
2.2.3.1	PCI Express* Port	26
2.2.3.2	PCI Express* Lanes Connection	26
2.3	Direct Media Interface (DMI)	27
2.3.1	DMI Error Flow	28
2.3.2	DMI Link Down	28
2.4	Processor Graphics	28
2.5	Processor Graphics Controller (GT)	29
2.5.1	3D and Video Engines for Graphics Processing	30
2.5.2	Multi Graphics Controllers Multi-Monitor Support	32
2.6	Digital Display Interface (DDI)	32
2.7	Intel® Flexible Display Interface (Intel® FDI)	38
2.8	Platform Environmental Control Interface (PECI)	38
2.8.1	PECI Bus Architecture	38
3	Technologies	41
3.1	Intel® Virtualization Technology (Intel® VT)	41
3.1.1	Intel® VT-x Objectives	41
3.1.2	Intel® VT-x Features	42
3.1.3	Intel® VT-d Objectives	43
3.1.4	Intel® VT-d Features	45
3.2	Intel® Trusted Execution Technology (Intel® TXT)	46
3.3	Intel® Hyper-Threading Technology (Intel® HT Technology)	47
3.4	Intel® Turbo Boost Technology 2.0	47
3.4.1	Intel® Turbo Boost Technology 2.0 Frequency	47
3.5	Intel® Advanced Vector Extensions 2.0 (Intel® AVX2)	48



3.5.1	AVX Disclaimer	48
3.6	Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)	48
3.6.1	PCLMULQDQ Instruction	49
3.6.2	Intel® Secure Key	49
3.7	Intel® Transactional Synchronization Extensions - New Instructions (Intel® TSX-NI)	49
3.8	Intel® 64 Architecture x2APIC	49
3.9	Power Aware Interrupt Routing (PAIR)	50
3.10	Execute Disable Bit	51
3.11	Supervisor Mode Execution Protection (SMEP)	51
4	Power Management	53
4.1	Advanced Configuration and Power Interface (ACPI) States Supported	54
4.2	Processor Core Power Management	55
4.2.1	Enhanced Intel SpeedStep® Technology Key Features	56
4.2.2	Low-Power Idle States	56
4.2.3	Requesting Low-Power Idle States	58
4.2.4	Core C-State Rules	58
4.2.5	Package C-States	60
4.2.6	Package C-States and Display Resolution	63
4.3	Integrated Memory Controller (IMC) Power Management	64
4.3.1	Disabling Unused System Memory Outputs	64
4.3.2	DRAM Power Management and Initialization	65
4.3.2.1	Initialization Role of CKE	66
4.3.2.2	Conditional Self-Refresh	66
4.3.2.3	Dynamic Power-Down	66
4.3.2.4	DRAM I/O Power Management	66
4.3.3	DRAM Running Average Power Limitation (RAPL)	67
4.3.4	DDR Electrical Power Gating (EPG)	67
4.4	PCI Express* Power Management	67
4.5	Direct Media Interface (DMI) Power Management	67
4.6	Graphics Power Management	67
4.6.1	Intel® Rapid Memory Power Management (Intel® RMPM)	67
4.6.2	Graphics Render C-State	67
4.6.3	Intel® Graphics Dynamic Frequency	68
5	Thermal Specifications Map	69
5.1	Processor Thermal Profiles	71
5.1.1	Processor (80W) Thermal Profile	72
5.1.2	Processor (65W) Thermal Profile	74
5.1.3	Processor (35W) Thermal Profile	75
5.2	Thermal Metrology	76
5.3	Fan Speed Control Scheme with Digital Thermal Sensor (DTS) 1.1	77
5.4	Fan Speed Control Scheme with Digital Thermal Sensor (DTS) 2.0	79
5.5	Product Family Thermal Specifications	80
5.5.1	Performance Targets	80
5.6	Processor Temperature	81
5.7	Adaptive Thermal Monitor	81
5.7.1	Frequency Control	82
5.7.2	Clock Modulation	83
5.7.3	Immediate Transition to Combined TM1 and TM2	83
5.7.4	Critical Temperature Flag	83
5.7.5	PROCHOT# Signal	84
5.8	THERMTRIP# Signal	85
5.9	Digital Thermal Sensor	85
5.9.1	Digital Thermal Sensor Accuracy (Taccuracy)	86
5.10	Intel® Turbo Boost Technology Thermal Considerations	86
5.11	Intel® Turbo Boost Technology Power Control and Reporting	86



5.11.1	Package Power Control.....	87
6	Signal Description	89
6.1	System Memory Interface Signals.....	89
6.2	Memory Reference and Compensation.....	91
6.3	Reset and Miscellaneous Signals.....	92
6.4	PCI Express*-Based Interface Signals	93
6.5	Display Interface Signals.....	93
6.6	Direct Media Interface (DMI).....	93
6.7	Phase Locked Loop (PLL) Signals	94
6.8	Testability Signals	94
6.9	Error and Thermal Protection Signals	95
6.10	Power Sequencing Signals	95
6.11	Processor Power Signals.....	96
6.12	Sense Signals	96
6.13	Ground and Non-Critical to Function (NCTF) Signals.....	96
6.14	Processor Internal Pull-Up / Pull-Down Terminations	97
7	Electrical Specifications	99
7.1	Integrated Voltage Regulator	99
7.2	Power and Ground Lands.....	99
7.3	VCC Voltage Identification (VID).....	99
7.4	Reserved or Unused Signals.....	106
7.5	Signal Groups	107
7.6	Test Access Port (TAP) Connection.....	109
7.7	DC Specifications	109
7.8	Voltage and Current Specifications.....	110
7.8.1	PECI DC Characteristics.....	115
7.8.2	Input Device Hysteresis	116
7.9	AC Specifications.....	116
7.10	Processor AC Timing Waveforms	126
7.11	Signal Quality	130
7.12	Overshoot / Undershoot Guidelines.....	131
7.12.1	VCC Overshoot Specification	131
7.12.2	Overshoot / Undershoot Magnitude.....	131
7.12.3	Overshoot / Undershoot Pulse Duration	131
8	Package Mechanical Specifications	133
8.1	Package Mechanical Drawing.....	133
8.2	Processor Component Keep-Out Zone	133
8.3	Package Loading Specifications	134
8.4	Package Handling Guidelines.....	134
8.5	Package Insertion Specifications.....	134
8.6	Processor Mass Specification	134
8.7	Processor Materials.....	135
8.8	Processor Markings.....	135
8.9	Processor Land Coordinates	135
8.10	Processor Storage Specifications	136
9	Processor Ball and Signal Information.....	139
10	DDR Data Swizzling	149



Figures

1-1	Platform Block Diagram	12
2-1	Intel® Flex Memory Technology Operations	22
2-2	PCI Express* Related Register Structures in the Processor	26
2-3	PCI Express* Typical Operation 16 Lanes Mapping	27
2-4	Processor Graphics Controller Unit Block Diagram	29
2-5	Processor Display Architecture	33
2-6	DisplayPort* Overview	34
2-7	HDMI* Overview	35
3-1	Device to Domain Mapping Structures	44
4-1	Processor Power States	53
4-2	Idle Power Management Breakdown of the Processor Cores	57
4-3	Thread and Core C-State Entry and Exit	57
4-4	Package C-State Entry and Exit	61
5-1	Thermal Test Vehicle Thermal Profile for Processor (80W)	72
5-2	Thermal Test Vehicle Thermal Profile for Processor (65W)	74
5-3	Thermal Test Vehicle Thermal Profile for Processor (35W)	75
5-4	Thermal Test Vehicle (TTV) Case Temperature (TCASE) Measurement Location	77
5-5	Digital Thermal Sensor (DTS) 1.1 Definition Points	78
5-6	Digital Thermal Sensor (DTS) Thermal Profile Definition	80
5-7	Package Power Control	88
7-1	Input Device Hysteresis	116
7-2	Differential Clock – Differential Measurements	126
7-3	Differential Clock – Single-Ended Measurements	127
7-4	DDR3L Command / Control and Clock Timing Waveform	127
7-5	DDR3L Data Setup and Hold Timing Waveform	128
7-6	DDR3L Clock to DQS Skew Timing Waveform	128
7-7	PCI Express* Receiver Eye Margins	129
7-8	TAP Valid Delay Timing Waveform	129
7-9	Test Reset (TRST#), Async Input, and PROCHOT# Output Timing Waveform	130
7-10	THERMTRIP# Power Down Sequence	130
7-11	Maximum Acceptable Overshoot / Undershoot Waveform	132
8-1	Processor Package Assembly Sketch	133
8-2	Processor Loading Specifications	134
8-3	Processor Top-Side Markings	135
8-4	Processor Package Land Coordinates	136



Tables

1-1	Terminology	14
1-2	Related Documents	18
2-1	Supported UDIMM Module Configurations	20
2-2	DDR3L System Memory Timing Support	21
2-3	PCI Express* Supported Configurations	24
2-4	Processor Supported Audio Formats over HDMI* and DisplayPort*	36
2-5	Valid Three Display Configurations through the Processor	36
2-6	DisplayPort and embedded DisplayPort* Resolutions for 1, 2, 4 Lanes – Link Data Rate of RBR, HBR, and HBR237	
4-1	System States	54
4-2	Processor Core / Package State Support	54
4-3	Integrated Memory Controller States	54
4-4	PCI Express* Link States	54
4-5	Direct Media Interface (DMI) States	55
4-6	G, S, and C Interface State Combinations	55
4-7	D, S, and C Interface State Combination	55
4-8	Coordination of Thread Power States at the Core Level	58
4-9	Coordination of Core Power States at the Package Level	61
4-10	Deepest Package C-State Available	63
5-1	Processor Thermal Specifications	70
5-2	Package Turbo Parameters	70
5-3	Thermal Test Vehicle Thermal Profile for Processor (80W)	72
5-4	Thermal Test Vehicle Thermal Profile for Processor (65W)	74
5-5	Thermal Test Vehicle Thermal Profile for Processor (35W)	76
5-6	Digital Thermal Sensor (DTS) 1.1 Thermal Solution Performance Above TCONTROL....	78
5-7	Thermal Margin Slope	80
5-8	Boundary Conditions, Performance Targets, and TCASE Specifications	81
5-9	Heatsink Information	81
6-1	Signal Description Buffer Types	89
6-2	Memory Channel A Signals	89
6-3	Memory Channel B Signals	90
6-4	Memory Reference and Compensation	91
6-5	Reset and Miscellaneous Signals	92
6-6	PCI Express* Graphics Interface Signals	93
6-7	Display Interface Signals	93
6-8	Direct Media Interface (DMI) – Processor to PCH Serial Interface	93
6-9	Phase Locked Loop (PLL) Signals	94
6-10	Testability Signals	94
6-11	Error and Thermal Protection Signals	95
6-12	Power Sequencing Signals	95
6-13	Processor Power Signals	96
6-14	Sense Signals	96
6-15	Ground and Non-Critical to Function (NCTF) Signals	96
6-16	Processor Internal Pull-Up / Pull-Down Terminations	97
7-1	VR 12.5 Voltage Identification	100
7-2	Signal Groups	107
7-3	Processor Core Active and Idle Mode DC Voltage and Current Specifications	110
7-4	Memory Controller (VDDQ) Supply DC Voltage and Current Specifications	111
7-5	VCCIO_OUT, VCOMP_OUT, and VCCIO_TERM	111



7-6	DDR3L Signal Group DC Specifications.....	112
7-7	Digital Display Interface Group DC Specifications	113
7-8	Embedded DisplayPort* (eDP) Group DC Specifications.....	113
7-9	CMOS Signal Group DC Specifications.....	114
7-10	GTL Signal Group and Open Drain Signal Group DC Specifications	114
7-11	PCI Express* DC Specifications.....	114
7-12	PECI DC Electrical Limits.....	115
7-13	Differential Clocks (SSC On).....	117
7-14	Differential Clocks (SSC Off).....	117
7-15	Processor Clock Jitter Specifications (Cycle-Cycle).....	118
7-16	System Reference Clock AC Specifications.....	118
7-17	DDR3L Electrical Characteristics and AC Timings at 1600 MT/s.....	119
7-18	DDR3L Electrical Characteristics and AC Timings at 1866 MT/s.....	120
7-19	PCI Express* AC Specification	120
7-20	PCI Express* Gen 3 Additional Specifications.....	122
7-21	Embedded DisplayPort* (eDP) AC Specifications	123
7-22	Display Port AC Specifications.....	123
7-23	HDMI* AC Specifications.....	124
7-24	Miscellaneous AC Specification.....	125
7-25	Testability Signal Group AC Specifications	125
7-26	SVID Signal Group AC Specifications	125
7-27	Processor Overshoot / Undershoot Specifications	132
8-1	Package Handling Guidelines	134
8-2	Processor Materials	135
8-3	Processor Storage Specifications.....	137
9-1	Processor Ball List by Signal Name	140



Revision History

Revision Number	Description	Revision Date
001	• Initial release	June 2015
002	• Updates to Thermal Specifications Chapter	June 2015

§





1 Introduction

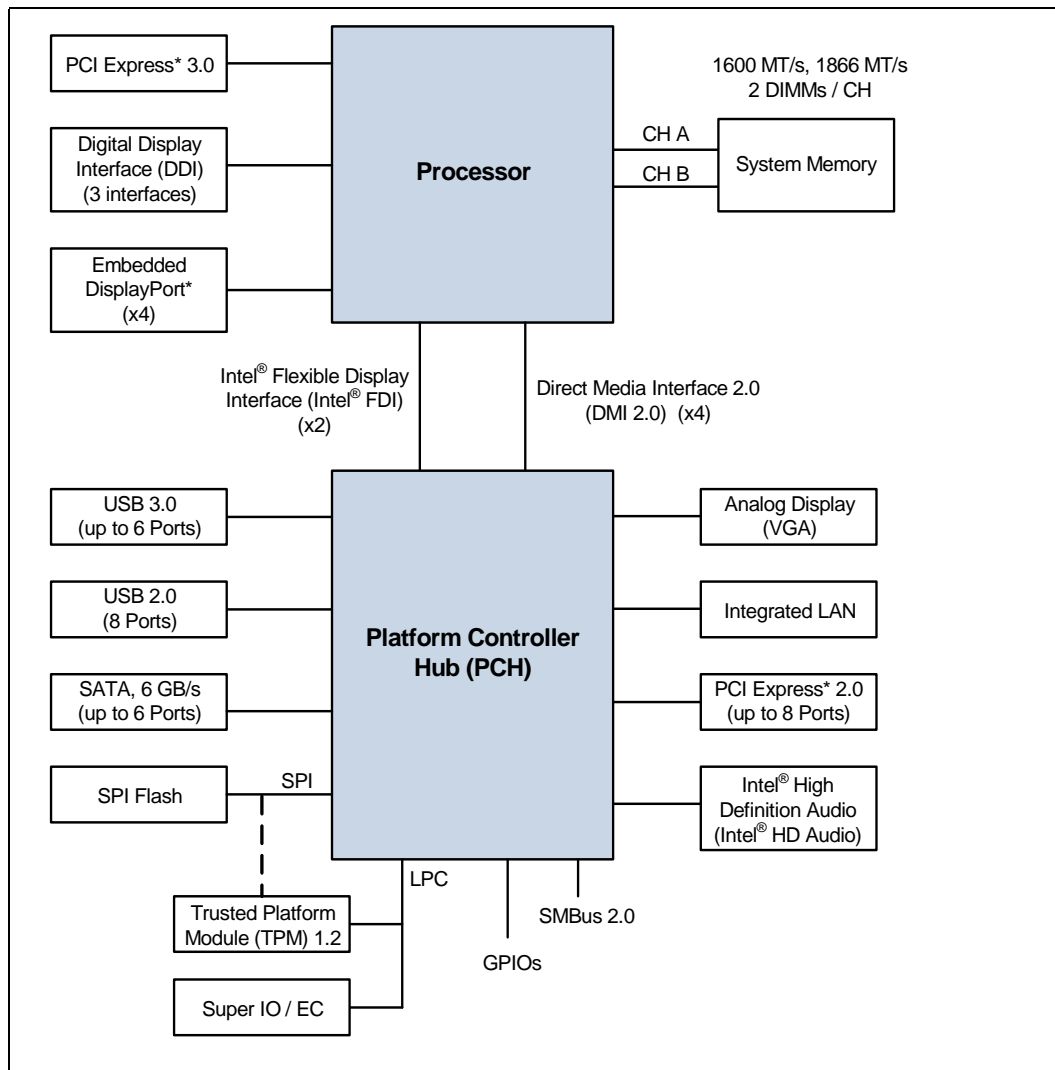
The Intel® Xeon® processor E3-1200 v4 product family (LGA -1150 socket H3) processors are 64-bit, multi-core processors built on 14-nanometer process technology.

The processors are designed for a two-chip platform consisting of a processor and Platform Controller Hub (PCH). The processors are designed to be used with the Intel® C226 Series Chipset Family Platform Controller Hub (PCH). These processors are “drop-in” compatible with the Desktop 4th Generation Intel® Core™ processor family and Intel® Xeon® E3-1200 v3 processor platforms. See the following figure for an example platform block diagram.

Note: Throughout this document, the Intel® Xeon® processor E3-1200 v4 product family may be referred to simply as “processor”.

Note: Throughout this document, the Intel® C226 Series Chipset Family Platform Controller Hub (PCH) may be referred to simply as “PCH”.

Figure 1-1. Platform Block Diagram



1.1 Supported Technologies

- Intel® Virtualization Technology (Intel® VT)
- Intel® Server Platform Services
- Intel® Trusted Execution Technology (Intel® TXT)
- Intel® Streaming SIMD Extensions 4.2 (Intel® SSE4.2)
- Intel® Hyper-Threading Technology (Intel® HT Technology)
- Intel® 64 Architecture
- Execute Disable Bit
- Intel® Turbo Boost Technology 2.0
- Intel® Advanced Vector Extensions 2.0 (Intel® AVX2)



- Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)
- Intel® PCLMULQDQ Instruction
- Intel® Device Protection Technology with Intel® Secure Key
- Intel® Transactional Synchronization Extensions - New Instructions (Intel® TSX-NI)
- PAIR – Power Aware Interrupt Routing
- SMEP – Supervisor Mode Execution Protection
- SMAP - Supervisor Mode Access Protection
- Intel® Device Protection Technology with Boot Guard
- DRAM Bit-Error Recovery (DBER)
- Enhanced Intel Speedstep® Technology

Note: The availability of the features may vary between processor SKUs.

1.2 Interfaces

The processor supports the following interfaces:

- DDR3L
- Direct Media Interface (DMI)
- Digital Display Interface (DDI)
- PCI Express*

1.3 Power Management Support

1.3.1 Processor Core

- Full support of ACPI C-states as implemented by the following processor C-states:
 - C0, C1, C1E, C3, C6
- Enhanced Intel SpeedStep® Technology

1.3.2 System

- S0, S3, S4, S5

1.3.3 Memory Controller

- Conditional self-refresh
- Dynamic power-down

1.3.4 PCI Express*

- L0s and L1 ASPM power management capability

1.3.5 DMI

- L0s and L1 ASPM power management capability



1.3.6 Processor Graphics Controller

- Intel® Rapid Memory Power Management (Intel® RMPM)
- Intel® Smart 2D Display Technology (Intel® S2DDT)
- Graphics Render C-state (RC6)
- Intel® Seamless Display Refresh Rate Switching with eDP port
- Intel® Display Power Saving Technology (Intel® DPST)

1.4 Thermal Management Support

- Digital Thermal Sensor
- 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

1.5 Package Support

The processor socket type (H3) is noted as LGA 1150. The package is a 37.5 x 37.5 mm Flip Chip Land Grid Array (FCLGA 1150).

1.6 Processor Testability

The processor includes boundary-scan for board and system level testability.

1.7 Terminology

Table 1-1. Terminology (Sheet 1 of 4)

Term	Description
APD	Active Power-down
B/D/F	Bus/Device/Function
BGA	Ball Grid Array
BLC	Backlight Compensation
BLT	Block Level Transfer
BPP	Bits per pixel
CKE	Clock Enable
CLTM	Closed Loop Thermal Management
DDI	Digital Display Interface



Table 1-1. Terminology (Sheet 2 of 4)

Term	Description
DDR3	Third-generation Double Data Rate SDRAM memory technology
DLL	Delay-Locked Loop
DMA	Direct Memory Access
DMI	Direct Media Interface
DP	DisplayPort*
DTS	Digital Thermal Sensor
DVI*	Digital Visual Interface. DVI* is the interface specified by the DDWG (Digital Display Working Group)
EC	Embedded Controller
ECC	Error Correction Code
eDP	Embedded Display Port
eDRAM	On Package Embedded DRAM
EPG	Electrical Power Gating
EU	Execution Unit
FMA	Floating-point fused Multiply Add instructions
FSC	Fan Speed Control
HDCP	High-bandwidth Digital Content Protection
HDMI*	High Definition Multimedia Interface
HFM	High Frequency Mode
iDCT	Inverse Discrete
IHS	Integrated Heat Spreader
GFX	Graphics
GSA	Graphics in System Agent
GUI	Graphical User 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® FDI	Intel Flexible Display Interface
Intel® TSX-NI	Intel Transactional Synchronization Extensions - New Instructions
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.
IOV	I/O Virtualization
ISI	Inter-Symbol Interference
ITPM	Integrated Trusted Platform Module
LCD	Liquid Crystal Display
LFM	Low Frequency Mode. LFM is Pn in the P-state table. It can be read at MSR CEh [47:40].
LFP	Local Flat Panel

Table 1-1. Terminology (Sheet 3 of 4)

Term	Description
MCP	Multi-Chip Package
MFM	Minimum Frequency Mode. MFM is the minimum ratio supported by the processor and can be read from MSR CEh [55:48].
MLE	Measured Launched Environment
MLC	Mid-Level Cache
MSI	Message Signaled Interrupt
MSL	Moisture Sensitive Labeling
MSR	Model Specific Registers
NCTF	Non-Critical to Function. NCTF locations are typically redundant ground or non-critical reserved, so the loss of the solder joint continuity at end of life conditions will not affect the overall product functionality.
ODT	On-Die Termination
OLTM	Open Loop Thermal Management
PCG	Platform Compatibility Guide (PCG) (previously known as FMB) provides a design target for meeting all planned processor frequency requirements.
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.
PECI	The Platform Environment Control Interface (PECI) is a one-wire interface that provides a communication channel between Intel processor and chipset components to external monitoring devices.
Ψ_{ca}	Case-to-ambient thermal characterization parameter (psi). A measure of thermal solution performance using total package power. Defined as $(T_{CASE} - T_{LA}) / \text{Total Package Power}$. The heat source should always be specified for Y measurements.
PEG	PCI Express* Graphics. External Graphics using PCI Express* Architecture. It is a high-speed serial interface where configuration is software compatible with the existing PCI specifications.
PL1, PL2	Power Limit 1 and Power Limit 2
PPD	Pre-charge Power-down
Processor	The 64-bit multi-core component (package)
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 L3 cache.
Processor Graphics	Intel Processor Graphics
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 SO-DIMM.
SCI	System Control Interrupt. SCI is used in the ACPI protocol.
SF	Strips and Fans
SMM	System Management Mode
SMX	Safer Mode Extensions
Storage Conditions	A non-operational state. The processor may be installed in a platform, in a tray, or loose. Processors may be sealed in packaging 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 must be handled in accordance with moisture sensitivity labeling (MSL) as indicated on the packaging material.
SVID	Serial Voltage Identification
TAC	Thermal Averaging Constant
TAP	Test Access Point


Table 1-1. Terminology (Sheet 4 of 4)

Term	Description
T_{CASE}	The case temperature of the processor, measured at the geometric center of the top-side of the TTV IHS.
TCC	Thermal Control Circuit
$T_{CONTROL}$	$T_{CONTROL}$ is a static value that is below the TCC activation temperature and used as a trigger point for fan speed control. When $DTS > T_{CONTROL}$, the processor must comply to the TTV thermal profile.
TDP	Thermal Design Power: Thermal solution should be designed to dissipate this target power level. TDP is not the maximum power that the processor can dissipate.
TLB	Translation Look-aside Buffer
TTV	Thermal Test Vehicle. A mechanically equivalent package that contains a resistive heater in the die to evaluate thermal solutions.
TM	Thermal Monitor. A power reduction feature designed to decrease temperature after the processor has reached its maximum operating temperature.
V_{CC}	Processor core power supply
V_{DDQ}	DDR3L power supply.
VF	Vertex Fetch
VID	Voltage Identification
VS	Vertex Shader
VLD	Variable Length Decoding
VMM	Virtual Machine Monitor
VR	Voltage Regulator
V_{SS}	Processor ground
x1	Refers to a Link or Port with one Physical Lane
x2	Refers to a Link or Port with two Physical Lanes
x4	Refers to a Link or Port with four Physical Lanes
x8	Refers to a Link or Port with eight Physical Lanes
x16	Refers to a Link or Port with sixteen Physical Lanes

1.8 Related Documents

Table 1-2. Related Documents (Sheet 1 of 2)

Document	Document Number / Location
Intel® Xeon® Processor E3-1200 v4 Product Family Datasheet Volume 2 of 2	332375
Intel® 9 Series Chipset Family Platform Controller Hub (PCH) Datasheet	330550
LGA1150 Socket Application Guide	328999
Advanced Configuration and Power Interface 3.0	http://www.acpi.info/
PCI Local Bus Specification 3.0	http://www.pcisig.com/specifications
PCI Express Base Specification, Revision 2.0	http://www.pcisig.com




Table 1-2. Related Documents (Sheet 2 of 2)

Document	Document Number / Location
DDR3 SDRAM Specification	http://www.jedec.org
DisplayPort* Specification	http://www.vesa.org
Intel® 64 and IA-32 Architectures Software Developer's Manuals	http://www.intel.com/products/processor/manuals/index.htm

§



2 Interfaces

2.1 System Memory Interface

- Two channels of DDR3L Unbuffered Dual In-Line Memory Modules (UDIMM) with a maximum of two DIMMs per channel.
- Single-channel and dual-channel memory organization modes
- Data burst length of eight for all memory organization modes
- Memory data transfer rates of 1600 MT/s, 1866 MT/s
- 64-bit wide channels
- DDR3L I/O voltage of 1.35 V for E3 Server
- The type of the DIMM modules supported by the processor is dependent on the PCH SKU in the target platform:
 - C226 PCH platforms support ECC and non-ECC UDIMMs
- Theoretical maximum memory bandwidth of:
 - 25.6 GB/s in dual-channel mode assuming DDR3L 1600 MT/s
 - 29.8 GB/s in dual-channel mode assuming DDR3L 1866 MT/s
- 2 Gb and 4 Gb DDR3L DRAM device technologies are supported
 - Using 4 Gb DRAM device technologies, the largest system memory capacity possible is 32 GB, assuming Dual Channel Mode with four x8 dual ranked DIMM memory configuration
- Up to 64 simultaneous open pages, 32 per channel (assuming 8 ranks of 8 bank devices)
- Processor on-die VREF generation for DDR DQ Read and Write as well as CMD/ADD
- Command launch modes of 1n/2n
- On-Die Termination (ODT)
- Asynchronous ODT
- Intel Fast Memory Access (Intel FMA):
 - Just-in-Time Command Scheduling
 - Command Overlap
 - Out-of-Order Scheduling
- eDRAM (Embedded DRAM):
 - 128 MB on Package
 - Up to 51 GB/s peak simultaneous read and write bandwidth added over system memory bandwidth



2.1.1 System Memory Technology Supported

The Integrated Memory Controller (IMC) supports DDR3L protocols with two independent, 64-bit wide channels each accessing one or two DIMMs. The type of memory supported by the processor is dependent on the PCH SKU in the target platform.

Note: The IMC supports a maximum of two DDR3L DIMMs per channel; thus, allowing up to four device ranks per channel.

Note: The support of DDR3L frequencies and number of DIMMs per channel is SKU dependent.

DDR3L Data Transfer Rates:

- 1600 MT/s (PC3-12800)
- 1866 MT/s, PC3-14900

Supported platforms UDIMM Modules:

- Raw Card B – Dual Ranked x8 unbuffered non-ECC
- Raw Card D – Single Ranked x8 unbuffered ECC
- Raw Card E – Dual Ranked x8 unbuffered ECC

Table 2-1. Supported UDIMM Module Configurations

Raw Card Version	DIMM Capacity	DRAM Device Technology	DRAM Organization	# of DRAM Devices	# of Physical Devices Ranks	# of Row / Col Address Bits	# of Banks Inside DRAM	Page Size
Unbuffered / Non-ECC Supported DIMM Module Configurations								
B	4 GB	2 Gb	256 M X 8	16	2	15/10	8	8K
	4 GB	4 Gb	512 M X 8	8	1	15/10	8	8K
	8 GB	4 Gb	512 M X 8	16	2	16/10	8	8K
Unbuffered / ECC Supported DIMM Module Configurations								
D	2 GB	2 Gb	256 M X 8	9	1	15/10	8	8K
	4 GB	2 Gb	256 M X 8	18	2	15/10	8	8K
	8 GB	4 Gb	512 M X 8	18	2	16/10	8	8K
E	4 GB	2 Gb	256 M X 8	18	2	15/10	8	8K
	8 GB	4 Gb	512 M X 8	18	2	16/10	8	8K

Note: DIMM module support is based on availability and is subject to change.

Note: System memory configurations are based on availability and are subject to change.



2.1.2 System Memory Timing Support

The IMC supports the following DDR3L 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
- CWL = CAS Write Latency
- Command Signal modes = 1N indicates a new command may be issued every clock and 2N indicates a new command may be issued every 2 clocks. Command launch mode programming depends on the transfer rate and memory configuration.

Table 2-2. DDR3L System Memory Timing Support

Segment	Transfer Rate (MT/s)	tCL (tCK)	tRCD (tCK)	tRP (tCK)	CWL (tCK)	DPC	CMD Mode
All segments	1600	10/11	10/11	10/11	8	1	1N/2N
						2	2N
	1866	13	12/13	12/13	9	1	1N/2N
						2	2N

Note: System memory timing support is based on availability and is subject to change.

2.1.3 System Memory Organization Modes

The Integrated Memory Controller (IMC) supports two memory organization modes – single-channel and dual-channel. Depending upon how the 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 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 symmetric and asymmetric zones. 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 must be greater or equal to channel B size.

Figure 2-1. Intel® Flex Memory Technology Operations



Dual-Channel Symmetric 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, the IMC operates completely in Dual-Channel Symmetric mode.

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



2.1.3.1 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 one or two DIMM connectors per channel. The usage of DIMM modules with different latencies is allowed, but in that case, the worst latency (among two channels) will be used. For dual-channel modes, both channels must have a DIMM connector populated and for single-channel mode only a single-channel may have one or both DIMM connectors populated.

Note: In a two-DIMM Per Channel (2DPC) layout memory configuration, the furthest DIMM from the processor of any given channel must always be populated first.

2.1.3.2 Intel® Fast Memory Access (Intel® FMA) Technology Enhancements

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.

2.1.3.3 Data Scrambling

The memory controller incorporates a DDR3L Data Scrambling feature to minimize the impact of excessive di/dt on the platform DDR3L 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 memory controller uses a data scrambling feature to create pseudo-random patterns on the DDR3L data bus to reduce the impact of any excessive di/dt.



2.2 PCI Express* Interface

This section describes the PCI Express* interface capabilities of the processor. See the *PCI Express Base Specification 3.0* for details on PCI Express*.

2.2.1 PCI Express* Support

The PCI Express* lanes (PEG[15:0] TX and RX) are fully-compliant to the *PCI Express Base Specification, Revision 3.0*.

The Intel® Xeon® processor with the following PCH SKU supports the configurations shown in the following table.

Table 2-3. PCI Express* Supported Configurations

Configuration	PCH C226
1x8, 2x4	GFX, I/O
2x8	GFX, I/O, Dual x8 GFX
1x16	GFX, I/O

- The port may negotiate down to narrower widths.
 - Support for x16/x8/x4/x2/x1 widths for a single PCI Express* mode.
- 2.5 GT/s, 5.0 GT/s and 8 GT/s PCI Express* bit rates are supported.
- Gen1 Raw bit-rate on the data pins of 2.5 GT/s, resulting in a real bandwidth per pair of 250 MB/s given the 8b/10b encoding used to transmit data across this interface. This also does not account for packet overhead and link maintenance. Maximum theoretical bandwidth on the interface of 4 GB/s in each direction simultaneously, for an aggregate of 8 GB/s when x16 Gen 1.
- Gen 2 Raw bit-rate on the data pins of 5.0 GT/s, resulting in a real bandwidth per pair of 500 MB/s given the 8b/10b encoding used to transmit data across this interface. This also does not account for packet overhead and link maintenance. Maximum theoretical bandwidth on the interface of 8 GB/s in each direction simultaneously, for an aggregate of 16 GB/s when x16 Gen 2.
- Gen 3 raw bit-rate on the data pins of 8.0 GT/s, resulting in a real bandwidth per pair of 984 MB/s using 128b/130b encoding to transmit data across this interface. This also does not account for packet overhead and link maintenance. Maximum theoretical bandwidth on the interface of 16 GB/s in each direction simultaneously, for an aggregate of 32 GB/s when x16 Gen 3.
- Hierarchical PCI-compliant configuration mechanism for downstream devices.
- Traditional PCI style traffic (asynchronous snooped, PCI ordering).
- PCI Express* extended configuration space. The first 256 bytes of configuration space aliases directly to the PCI Compatibility configuration space. The remaining portion of the fixed 4-KB block of memory-mapped space above that (starting at 100h) is known as extended configuration space.
- PCI Express* Enhanced Access Mechanism. Accessing the device configuration space in a flat memory mapped fashion.
- Automatic discovery, negotiation, and training of link out of reset.
- Traditional AGP style traffic (asynchronous non-snooped, PCI-X* Relaxed ordering).



- Peer segment destination posted write traffic (no peer-to-peer read traffic) in Virtual Channel 0: DMI -> PCI Express* Port 0.
- 64-bit downstream address format, but the processor never generates an address above 64 GB (Bits 63:36 will always be zeros).
- 64-bit upstream address format, but the processor responds to upstream read transactions to addresses above 64 GB (addresses where any of Bits 63:36 are nonzero) with an Unsupported Request response. Upstream write transactions to addresses above 64 GB will be dropped.
- Re-issues Configuration cycles that have been previously completed with the Configuration Retry status.
- PCI Express* reference clock is 100-MHz differential clock.
- Power Management Event (PME) functions.
- Dynamic width capability.
- Message Signaled Interrupt (MSI and MSI-X) messages.
- Polarity inversion

Note: The processor does not support PCI Express* Hot-Plug.

2.2.2 PCI Express* Architecture

Compatibility with the PCI addressing model is maintained to ensure that all existing applications and drivers operate unchanged.

The PCI Express* configuration uses standard mechanisms as defined in the PCI Plug-and-Play specification. The processor PCI Express* ports support Gen 3. At 8 GT/s, Gen 3 operation results in twice as much bandwidth per lane as compared to Gen 2 operation. The 16 lanes PEG can operate at 2.5 GT/s, 5 GT/s, or 8 GT/s.


Gen 3 PCI Express* uses a 128b/130b encoding that is about 23% more efficient than the 8b/10b encoding used in Gen 1 and Gen 2.

The PCI Express* architecture is specified in three layers – Transaction Layer, Data Link Layer, and Physical Layer. See the *PCI Express Base Specification 3.0* for details of PCI Express* architecture.

2.2.3 PCI Express* Configuration Mechanism

The PCI Express* (external graphics) link is mapped through a PCI-to-PCI bridge structure.

Figure 2-2. PCI Express* Related Register Structures in the Processor



PCI Express* extends the configuration space to 4096 bytes per-device/function, as compared to 256 bytes allowed by the conventional PCI specification. PCI Express* configuration space is divided into a PCI-compatible region (that consists of the first 256 bytes of a logical device's configuration space) and an extended PCI Express* region (that consists of the remaining configuration space). The PCI-compatible region can be accessed using either the mechanisms defined in the PCI specification or using the enhanced PCI Express* configuration access mechanism described in the PCI Express* Enhanced Configuration Mechanism section.

The PCI Express* Host Bridge is required to translate the memory-mapped PCI Express* configuration space accesses from the host processor to PCI Express* configuration cycles. To maintain compatibility with PCI configuration addressing mechanisms, it is recommended that system software access the enhanced configuration space using 32-bit operations (32-bit aligned) only. See the *PCI Express Base Specification* for details of both the PCI-compatible and PCI Express* Enhanced configuration mechanisms and transaction rules.

2.2.3.1 PCI Express* Port


The PCI Express* interface on the processor is a single, 16-lane (x16) port that can also be configured at narrower widths. The PCI Express* port is being designed to be compliant with the *PCI Express Base Specification*, Revision 3.0.

2.2.3.2 PCI Express* Lanes Connection

The following figure demonstrates the PCIe* lane mapping.



Figure 2-3. PCI Express* Typical Operation 16 Lanes Mapping



2.3 Direct Media Interface (DMI)

Direct Media Interface (DMI) connects the processor and the PCH. Next generation DMI2 is supported.

Note:

Only DMI x4 configuration is supported.

- DMI 2.0 support.
- Compliant to Direct Media Interface Second Generation (DMI2).
- Four lanes in each direction.
- 5 GT/s point-to-point DMI interface to PCH is supported.
- Raw bit-rate on the data pins of 5.0 GB/s, resulting in a real bandwidth per pair of 500 MB/s given the 8b/10b encoding used to transmit data across this interface. Does not account for packet overhead and link maintenance.
- Maximum theoretical bandwidth on interface of 2 GB/s in each direction simultaneously, for an aggregate of 4 GB/s when DMI x4.
- Shares 100-MHz PCI Express* reference clock.

- 64-bit downstream address format, but the processor never generates an address above 64 GB (Bits 63:36 will always be zeros).
- 64-bit upstream address format, but the processor responds to upstream read transactions to addresses above 64 GB (addresses where any of Bits 63:36 are nonzero) with an Unsupported Request response. Upstream write transactions to addresses above 64 GB will be dropped.
- Supports the following traffic types to or from the PCH:
 - DMI -> DRAM
 - DMI -> processor core (Virtual Legacy Wires (VLWs), Resetwarn, or MSIs only)
 - Processor core -> DMI
- APIC and MSI interrupt messaging support:
 - Message Signaled Interrupt (MSI and MSI-X) messages
- Downstream SMI, SCI and SERR error indication.
- Legacy support for ISA regime protocol (PHOLD/PHOLDA) required for parallel port DMA, floppy drive, and LPC bus masters.
- DC coupling – no capacitors between the processor and the PCH.
- Polarity inversion.
- PCH end-to-end lane reversal across the link.
- Supports Half Swing “low-power/low-voltage”.

2.3.1 DMI Error Flow

DMI can only generate SERR in response to errors, never SCI, SMI, MSI, PCI INT, or GPE. Any DMI related SERR activity is associated with Device 0.

2.3.2 DMI Link Down

The DMI link going down is a fatal, unrecoverable error. If the DMI data link goes to data link down, after the link was up, then the DMI link hangs the system by not allowing the link to retrain to prevent data corruption. This link behavior is controlled by the PCH.

Downstream transactions that had been successfully transmitted across the link prior to the link going down may be processed as normal. No completions from downstream, non-posted transactions are returned upstream over the DMI link after a link down event.

2.4 Processor Graphics

The processor graphics contains a generation 8.0 graphics core architecture. This enables substantial gains in performance and lower power consumption over previous generations. 48 Execution Units (per graphics core) are supported.

- Next Generation Intel Clear Video Technology HD Support is a collection of video playback and enhancement features that improve the end user’s viewing experience
 - Encode / transcode HD content
 - Playback of high definition content including Blu-ray Disc*
 - Superior image quality with sharper, more colorful images
 - Playback of Blu-ray* disc S3D content using HDMI (1.4a specification compliant with 3D)




- DirectX* Video Acceleration (DXVA) support for accelerating video processing
 - Full AVC/VC1/MPEG2 HW Decode
- Advanced Scheduler 2.0, 1.0, XPDM support
- Windows* 8.x, Windows* 7, OSX, Linux* operating system support
- DirectX* 11.1, DirectX* 11, DirectX* 10.1, DirectX* 10, DirectX* 9 support.
- OpenGL* 4.0, support

2.5 Processor Graphics Controller (GT)

The New Graphics Engine 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.

Figure 2-4. Processor Graphics Controller Unit Block Diagram





2.5.1 3D and Video Engines for Graphics Processing

The Gen 8.0 3D engine provides the following performance and power-management enhancements.

3D Pipeline

The 3D graphics pipeline architecture simultaneously operates on different primitives or on different portions of the same primitive. All the cores are fully programmable, increasing the versatility of the 3D Engine.

3D Engine Execution Units

- Supports 48 EUs. The EUs perform 128-bit wide execution per clock.
- Support SIMD8 instructions for vertex processing and SIMD16 instructions for pixel processing.

Vertex Fetch (VF) Stage

The VF stage executes 3DPRIMITIVE commands. Some enhancements have been included to better support legacy D3D APIs as well as SGI OpenGL*.

Vertex Shader (VS) Stage

The VS stage performs shading of vertices output by the VF function. The VS unit produces an output vertex reference for every input vertex reference received from the VF unit, in the order received.

Geometry Shader (GS) Stage

The GS stage receives inputs from the VS stage. Compiled application-provided GS programs, specifying an algorithm to convert the vertices of an input object into some output primitives. For example, a GS shader may convert lines of a line strip into polygons representing a corresponding segment of a blade of grass centered on the line. Or it could use adjacency information to detect silhouette edges of triangles and output polygons extruding out from the edges.

Clip Stage

The Clip stage performs general processing on incoming 3D objects. However, it also includes specialized logic to perform a Clip Test function on incoming objects. The Clip Test optimizes generalized 3D Clipping. The Clip unit examines the position of incoming vertices, and accepts/rejects 3D objects based on its Clip algorithm.

Strips and Fans (SF) Stage

The SF stage performs setup operations required to rasterize 3D objects. The outputs from the SF stage to the windower stage contain implementation-specific information required for the rasterization of objects and also supports clipping of primitives to some extent.

Windower / IZ (WIZ) Stage

The WIZ unit performs an early depth test, which removes failing pixels and eliminates unnecessary processing overhead.



The windower uses the parameters provided by the SF unit in the object-specific rasterization algorithms. The WIZ unit rasterizes objects into the corresponding set of pixels. The windower is also capable of performing dithering, whereby the illusion of a higher resolution when using low-bpp channels in color buffers is possible. Color dithering diffuses the sharp color bands seen on smooth-shaded objects.

Video Engine

The Video Engine handles the non-3D (media/video) applications. It includes support for VLD and MPEG2 decode in hardware.

2D Engine

The 2D Engine contains BLT (Block Level Transfer) functionality and an extensive set of 2D instructions. To take advantage of the 3D during engine's functionality, some BLT functions make use of the 3D renderer.

Processor Graphics VGA Registers

The 2D registers consists of original VGA registers and others to support graphics modes that have color depths, resolutions, and hardware acceleration features that go beyond the original VGA standard.

Logical 128-Bit Fixed BLT and 256 Fill Engine

This BLT engine accelerates the GUI of Microsoft Windows* operating systems. The 128-bit BLT engine provides hardware acceleration of block transfers of pixel data for many common Windows operations. The BLT engine can be used for the following:

- Move rectangular blocks of data between memory locations
- Data alignment
- To perform logical operations (raster ops)

The rectangular block of data does not change, as it is transferred between memory locations. The allowable memory transfers are between: cacheable system memory and frame buffer memory, frame buffer memory and frame buffer memory, and within system memory. Data to be transferred can consist of regions of memory, patterns, or solid color fills. A pattern is always 8 x 8 pixels wide and may be 8, 16, or 32 bits per pixel.

The BLT engine expands monochrome data into a color depth of 8, 16, or 32 bits. BLTs can be either opaque or transparent. Opaque transfers move the data specified to the destination. Transparent transfers compare destination color to source color and write according to the mode of transparency selected.

Data is horizontally and vertically aligned at the destination. If the destination for the BLT overlaps with the source memory location, the BLT engine specifies which area in memory to begin the BLT transfer. Hardware is included for all 256 raster operations (source, pattern, and destination) defined by Microsoft*, including transparent BLT.

The BLT engine has instructions to invoke BLT and stretch BLT operations, permitting software to set up instruction buffers and use batch processing. The BLT engine can perform hardware clipping during BLTs.



2.5.2 Multi Graphics Controllers Multi-Monitor Support

The processor supports simultaneous use of the Processor Graphics Controller (GT) and a x16 PCI Express* Graphics (PEG) device. The processor supports a maximum of 2 displays connected to the PEG card in parallel with up to 2 displays connected to the processor and PCH.

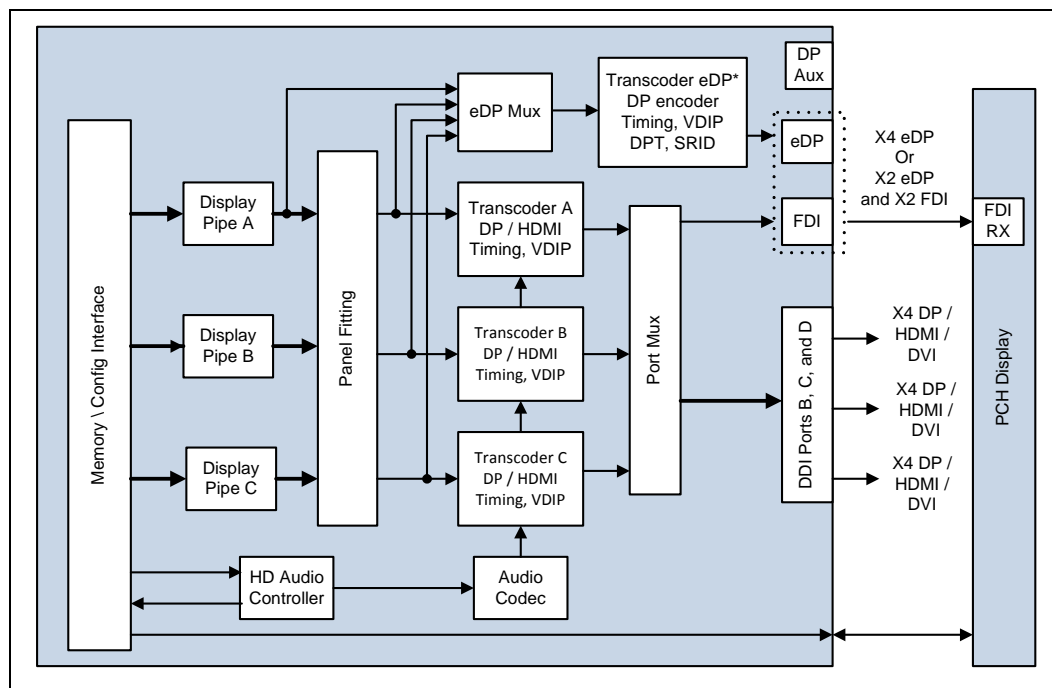
Note: When supporting Multi Graphics Multi Monitors, "drag and drop" between monitors and the 2x8PEG is not supported.

2.6 Digital Display Interface (DDI)

- The processor supports:
 - Three Digital Display (x4 DDI) interfaces that can be configured as DisplayPort*, HDMI*, or DVI. DisplayPort* can be configured to use 1, 2, or 4 lanes depending on the bandwidth requirements and link data rate of RBR (1.62 GT/s), HBR (2.7 GT/s) and HBR2 (5.4 GT/s). When configured as HDMI*, DDIx4 port can support 2.97 GT/s. In addition, Digital Port D (x4 DDI) interface can also be configured to carry embedded DisplayPort* (eDPx4). Built-in displays are only supported on Digital Port D.
 - One dedicated Intel FDI Port for legacy VGA support on the PCH.
- The HDMI* interface supports HDMI with 3D, 4K, Deep Color, and x.v.Color. The DisplayPort* interface supports the VESA DisplayPort* Standard Version 1, Revision 2.
- The processor supports High-bandwidth Digital Content Protection (HDCP) for high-definition content playback over digital interfaces.
- The processor also integrates dedicated a Mini HD audio controller to drive audio on integrated digital display interfaces, such as HDMI* and DisplayPort*. The HD audio controller on the PCH would continue to support down CODECs, and so on. The processor Mini HD audio controller supports two High-Definition Audio streams simultaneously on any of the three digital ports.
- The processor supports streaming any 3 independent and simultaneous display combination of DisplayPort*/HDMI*/DVI/eDP*/VGA monitors with the exception of 3 simultaneous display support of HDMI*/DVI. In the case of 3 simultaneous displays, two High Definition Audio streams over the digital display interfaces are supported.
- Each digital port is capable of driving resolutions up to 3840x2160 at 60 Hz through DisplayPort* and 4096x2304 at 24 Hz/2560x1600 at 60 Hz using HDMI*.
- DisplayPort* Aux CH, DDC channel, Panel power sequencing, and HPD are supported through the PCH.



Figure 2-5. Processor Display Architecture



Display is the presentation stage of graphics. This involves:

- Pulling rendered data from memory
- Converting raw data into pixels
- Blending surfaces into a frame
- Organizing pixels into frames
- Optionally scaling the image to the desired size
- Re-timing data for the intended target
- Formatting data according to the port output standard


DisplayPort*

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. DisplayPort* is also suitable for display connections between consumer electronics devices, such as high-definition optical disc players, set top boxes, and TV displays.

A DisplayPort* consists of a Main Link, 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 for the sink device.

The processor is designed in accordance with the VESA DisplayPort* Standard Version 1.2a. The processor supports VESA DisplayPort* PHY Compliance Test Specification 1.2a and VESA DisplayPort* Link Layer Compliance Test Specification 1.2a.

Figure 2-6. DisplayPort* Overview




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 audiovisual 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 with 3D, 4K, Deep Color, and x.v.Color.

Figure 2-7. HDMI* Overview


Digital Video Interface

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. To drive DVI-I through the back panel the VGA DDC signals are connected along with the digital data and clock signals from one of the Digital Ports. When a system has support for a DVI-I port, then either VGA or the DVI-D through a single DVI-I connector can be driven, but not both simultaneously.

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.

Embedded DisplayPort*

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

The eDP on the processor can be configured for 2 or 4 lanes.

The processor supports Embedded DisplayPort* (eDP*) Standard Version 1.2 and VESA Embedded DisplayPort* Standard Version 1.2.

Integrated Audio

- HDMI and display port interfaces carry audio along with video.
- Processor supports two DMA controllers to output two High Definition audio streams on two digital ports simultaneously.
- Supports only the internal HDMI and DP CODECs.

Table 2-4. 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, 8 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 44.1 kHz, 48 kHz, 88.2 kHz, 96 kHz, 176.4 kHz, and 192 kHz sampling rates.

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.
- Intel 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.

The digital ports on the processor can be configured to support DisplayPort*/HDMI/DVI. Digital port D can be configured as eDPx4 in addition to dedicated x2 port for Intel FDI for VGA. The following table shows examples of valid three display configurations through the processor.

Table 2-5. Valid Three Display Configurations through the Processor (Sheet 1 of 2)

Display 1	Display 2	Display 3	Maximum Resolution Display 1	Maximum Resolution Display 2	Maximum Resolution Display 3
HDMI	HDMI	DP	4096x2304 @ 24 Hz 2560x1600 @ 60 Hz		3840x2160 @ 60 Hz
DVI	DVI	DP	1920x1200 @ 60 Hz		3840x2160 @ 60 Hz
DP	DP	DP	3840x2160 @ 60 Hz		
VGA	DP	HDMI	1920x1200 @ 60 Hz	3840x2160 @ 60 Hz	4096x2304 @ 24 Hz 2560x1600 @ 60 Hz
eDP	DP	HDMI	3840x2160 @ 60 Hz	3840x2160 @ 60 Hz	4096x2304 @ 24 Hz 2560x1600 @ 60 Hz



Table 2-5. Valid Three Display Configurations through the Processor (Sheet 2 of 2)

Display 1	Display 2	Display 3	Maximum Resolution Display 1	Maximum Resolution Display 2	Maximum Resolution Display 3
eDP	DP	DP	3840x2160 @ 60 Hz	3840x2160 @ 60 Hz	
eDP	HDMI	HDMI	3840x2160 @ 60 Hz	4096x2304 @ 24 Hz 2560x1600 @ 60 Hz	

Notes:

- Requires support of 2 channel DDR3L 1600 MT/s / 1866 MT/s configuration for driving 3 simultaneous 3840x2160 @ 60 Hz display resolutions
- DP and eDP resolutions in the above table are supported for 4 lanes with link data rate HBR2.

The following table shows the DP/eDP resolutions supported for 1, 2, or 4 lanes depending on link data rate of RBR, HBR, and HBR2.

Table 2-6. DisplayPort and embedded DisplayPort* Resolutions for 1, 2, 4 Lanes – Link Data Rate of RBR, HBR, and HBR2

Link Data Rate	Lane Count		
	1	2	4
RBR	1064x600	1400x1050	2240x1400
HBR	1280x960	1920x1200	2880x1800
HBR2	1920x1200	2880x1800	3840x2160

Any 3 displays can be supported simultaneously using the following rules:

- Maximum of 2 HDMI
- Maximum of 2 DVIs
- Maximum of 1 HDMI and 1 DVI
- Any 3 DisplayPort
- One VGA
- One eDP

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 HDCP 1.4 for content protection over wired displays (HDMI*, DVI, and DisplayPort*).

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

2.7 Intel[®] Flexible Display Interface (Intel[®] FDI)

- The Intel Flexible Display Interface (Intel FDI) passes display data from the processor (source) to the PCH (sink) for display through a display interface on the PCH.
- Intel FDI supports 2 lanes at 2.7 GT/s fixed frequency. This can be configured to 1 or 2 lanes depending on the bandwidth requirements.
- Intel FDI supports 8 bits per color only.
- Side band sync pin (FDI_CSUNC).
- Side band interrupt pin (DISP_INT). This carries combined interrupt for HPDs of all the ports, AUX and I²C completion events, and so on.
- Intel FDI is not encrypted as it drives only VGA and content protection is not supported on VGA.

2.8 Platform Environmental Control Interface (PECI)

PECI is an Intel proprietary interface that provides a communication channel between Intel processors and external components like Super I/O (SIO) and Embedded Controllers (EC) to provide processor temperature, Turbo, 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.

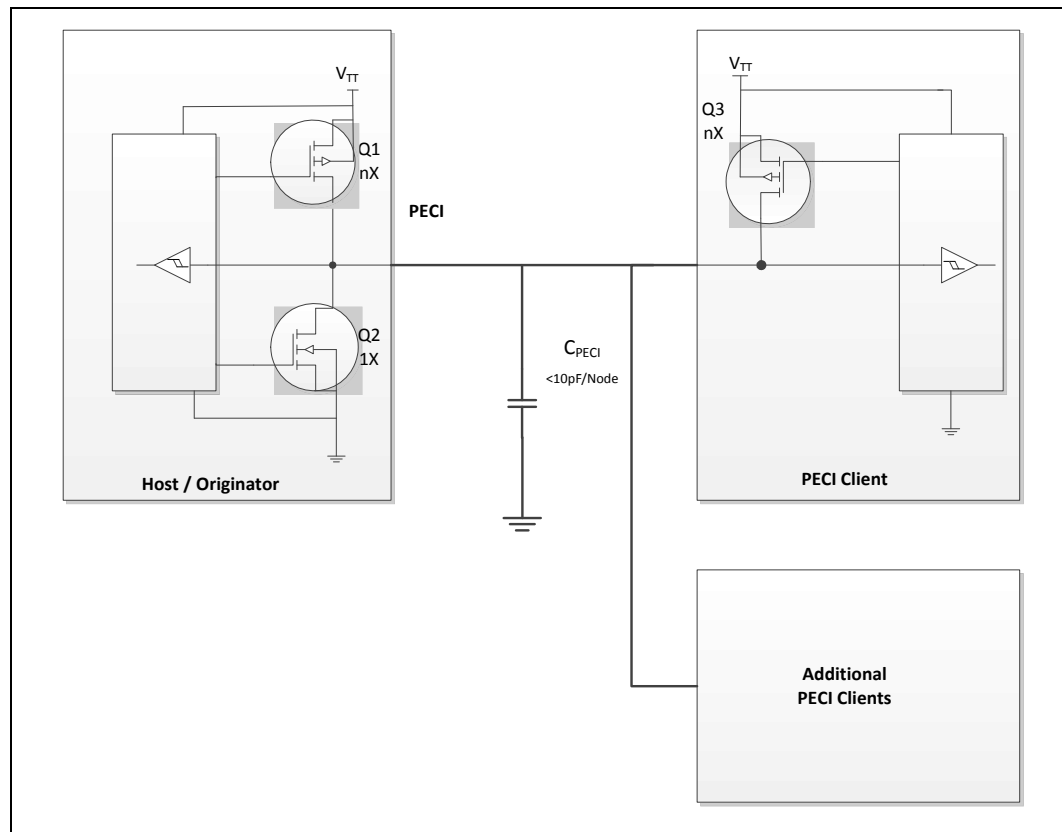
2.8.1 PECI Bus Architecture

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

The idle state on the bus is near zero.

The following figure demonstrates PECI design and connectivity. While the host/originator can be third party PECI host, and one of the PECI client is a processor PECI device.

Equation 2-1. PECI Host-Clients Connection Example



§





3 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/>

3.1 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) for IA-32, 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 and IA-32 Architectures Software Developer's Manual, Volume 3B* and is available at:

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

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

<http://www.intel.com/technology/virtualization/index.htm>

<https://sharedspaces.intel.com/sites/PCDC/SitePages/Ingredients/ingredient.aspx?ing=VT>

3.1.1 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 they 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.

3.1.2 Intel® VT-x Features

The processor supports the following 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.
- Extended Page Table Pointer (EPTP) 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 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 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.




3.1.3 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 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 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 (Intel VT-d Table) that contains Guest specific address translations.

Figure 3-1. 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://download.intel.com/technology/computing/vptech/Intel\(r\)_VT_for_Direct_IO.pdf](http://download.intel.com/technology/computing/vptech/Intel(r)_VT_for_Direct_IO.pdf)

3.1.4 Intel® VT-d Features

The processor supports the following Intel VT-d features:

- Memory controller and processor graphics comply with the Intel VT-d 1.2 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 4K page sizes
- 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 (Mem Wr to address FEEEx_xxxxh) not translated
 - Translation faults result in cycle forwarding to VBIOS region (byte enables masked for writes). Returned data may be bogus for internal agents; PEG/DMI interfaces return unsupported request status
- 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 IGD VT-d engine, are upgraded to support 4-level Intel VT-d tables (adjusted guest address width 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.



3.2 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 MSRs to:

- Enable a second SMM range
- Enable SMM code execution range checking
- Select whether SMM Save State is to be written to legacy SMRAM or to MSRs
- 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 must test the associated capability bit before attempting to access any of the above registers.



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

The processor supports Intel Hyper-Threading Technology (Intel HT Technology) that allows an execution 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 must be enabled using the BIOS and requires operating system support.

Intel recommends enabling Intel HT Technology with Microsoft Windows* 8 and Microsoft Windows* 7 and disabling Intel HT Technology using the BIOS for all previous versions of Windows* operating systems. For more information on Intel HT Technology, see <http://www.intel.com/technology/platform-technology/hyper-threading/>.

3.4 Intel® Turbo Boost Technology 2.0

The Intel Turbo Boost Technology 2.0 allows the processor core to opportunistically and automatically run faster than its rated operating frequency/render clock 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.

Maximum frequency is dependant on the SKU and number of active cores. No special hardware support is necessary for Intel Turbo Boost Technology 2.0. BIOS and the operating system can enable or disable Intel Turbo Boost Technology 2.0.

Compared with previous generation products, Intel Turbo Boost Technology 2.0 will increase the ratio of application power to TDP. 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.

3.4.1 Intel® Turbo Boost Technology 2.0 Frequency

The processor rated frequency assumes that all execution cores are running an application at the thermal design power (TDP). However, under typical operation, not all cores are active. Therefore, most applications are consuming less than the TDP at the rated frequency. To take advantage of the available thermal headroom, the active cores can increase their operating frequency.

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

- The number of cores operating in the C0 state.
- The estimated core current consumption.
- The estimated package prior and present power consumption.
- 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, see [Chapter 4, “Power Management”](#).

3.5 Intel® Advanced Vector Extensions 2.0 (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, see <http://www.intel.com/software/avx>

3.5.1 AVX Disclaimer

Intel® Advanced Vector Extensions (Intel® AVX)* are designed to achieve higher throughput for certain integer and floating point operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and 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 you should consult your system manufacturer for more information.

*** Intel® Advanced Vector Extensions refers to Intel® AVX, Intel® AVX2 or Intel® AVX-512. For more information on Intel® Turbo Boost Technology 2.0, visit <http://www.intel.com/go/turbo>**

3.6 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 a full hardware for supporting AES; offering security, high performance, and a great deal of flexibility.



3.6.1 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.

3.6.2 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, and so on.

3.7 Intel® Transactional Synchronization Extensions - New Instructions (Intel® TSX-NI)

Intel Transactional Synchronization Extensions - New Instructions (Intel TSX-NI). Intel TSX-NI provides a set of instruction set extensions that allow programmers to specify regions of code for transactional synchronization. Programmers can use these extensions to achieve the performance of fine-grain locking while actually programming using coarse-grain locks. Details on Intel TSX-NI may be found in *Intel® Architecture Instruction Set Extensions Programming Reference*.

3.8 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.
 - In x2APIC mode, APIC registers are accessed through Model Specific Register (MSR) interfaces. In this mode, the x2APIC architecture provides significantly increased processor addressability and some enhancements on interrupt delivery.
- 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.
- More efficient MSR interface to access APIC registers:
 - To enhance inter-processor and self-directed interrupt delivery as well as the ability to virtualize the local APIC, the APIC register set can be accessed only through MSR-based interfaces in x2APIC mode. The Memory Mapped IO (MMIO) interface used by xAPIC is not supported in x2APIC mode.
- 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 extendible for future Intel platform innovations.

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

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

3.9 Power Aware Interrupt Routing (PAIR)

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



3.10 Execute Disable Bit

The Execute Disable Bit allows memory to be marked as 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. See the *Intel® 64 and IA-32 Architectures Software Developer's Manuals* for more detailed information.

3.11 Supervisor Mode Execution Protection (SMEP)

The processor introduces a new 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. For more information, please refer to *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A* at: <http://www.intel.com/Assets/PDF/manual/253668.pdf>

§



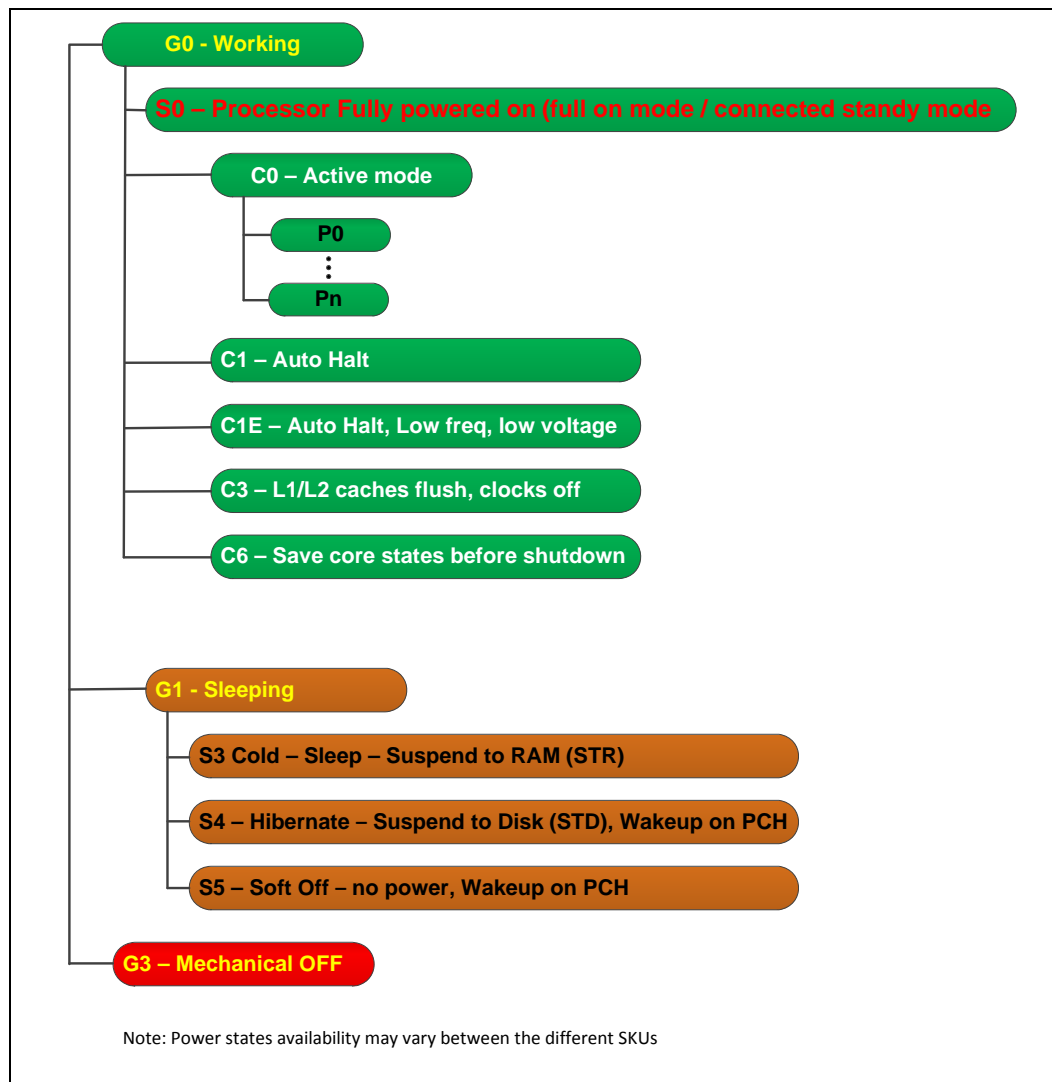


4 Power Management

This chapter provides information on the following power management topics:

- Advanced Configuration and Power Interface (ACPI) States
- Processor Core
- Integrated Memory Controller (IMC)
- PCI Express*
- Direct Media Interface (DMI)
- Processor Graphics Controller

Figure 4-1. Processor Power States





4.1 Advanced Configuration and Power Interface (ACPI) States Supported

This section describes the ACPI states supported by the processor.

Table 4-1. System States

State	Description
G0/S0	Full On Mode.
G1/S3-Cold	Suspend-to-RAM (STR). Context saved to memory (S3-Hot state is not supported by the processor).
G1/S4	Suspend-to-Disk (STD). All power lost (except wake up on PCH).
G2/S5	Soft off. All power lost (except wake up on PCH). Total reboot.
G3	Mechanical off. All power removed from system.

Table 4-2. Processor Core / Package State Support

State	Description
C0	Active mode, processor executing code.
C1	AutoHALT state.
C1E	AutoHALT state with lowest frequency and voltage operating point.
C3	Execution cores in C3 state flush their L1 instruction cache, L1 data cache, and L2 cache to the L3 shared cache. Clocks are shut off to each core.
C6	Execution cores in this state save their architectural state before removing core voltage.

Table 4-3. Integrated Memory Controller 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 4-4. PCI Express* Link States

State	Description
L0	Full on – Active transfer state.
L0s	First Active Power Management low power state – Low exit latency.
L1	Lowest Active Power Management – Longer exit latency.
L3	Lowest power state (power-off) – Longest exit latency.



Table 4-5. Direct Media Interface (DMI) States

State	Description
L0	Full on – Active transfer state.
L0s	First Active Power Management low power state – Low exit latency.
L1	Lowest Active Power Management – Longer exit latency.
L3	Lowest power state (power-off) – Longest exit latency.

Table 4-6. 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	C1/C1E	Auto-Halt	On	Auto-Halt
G0	S0	C3	Deep Sleep	On	Deep Sleep
G0	S0	C6	Deep Power-down	On	Deep Power-down
G1	S3	Power off		Off, except RTC	Suspend to RAM
G1	S4	Power off		Off, except RTC	Suspend to Disk
G2	S5	Power off		Off, except RTC	Soft Off
G3	NA	Power off		Power off	Hard off

Table 4-7. D, S, and C Interface State Combination

Graphics Adapter (D) State	Sleep (S) State	Package (C) State	Description
D0	S0	C0	Full On, Displaying.
D0	S0	C1/C1E	Auto-Halt, Displaying.
D0	S0	C3	Deep sleep, Displaying.
D0	S0	C6	Deep Power-down, Displaying.
D3	S0	Any	Not displaying.
D3	S3	N/A	Not displaying, Graphics Core is powered off.
D3	S4	N/A	Not displaying, suspend to disk.

4.2 Processor Core Power Management

While executing code, Enhanced Intel SpeedStep® Technology optimizes the processor’s frequency and core 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.



4.2.1 Enhanced Intel SpeedStep® Technology Key Features

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 MSRs. The voltage is optimized based on the selected frequency and the number of active processor cores.
 - Once the voltage is established, the PLL locks on to the target frequency.
 - All active processor cores share the same frequency and voltage. In a multi-core processor, the highest frequency P-state requested among all active 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.
- Because there is low transition latency between P-states, a significant number of transitions per-second are possible.

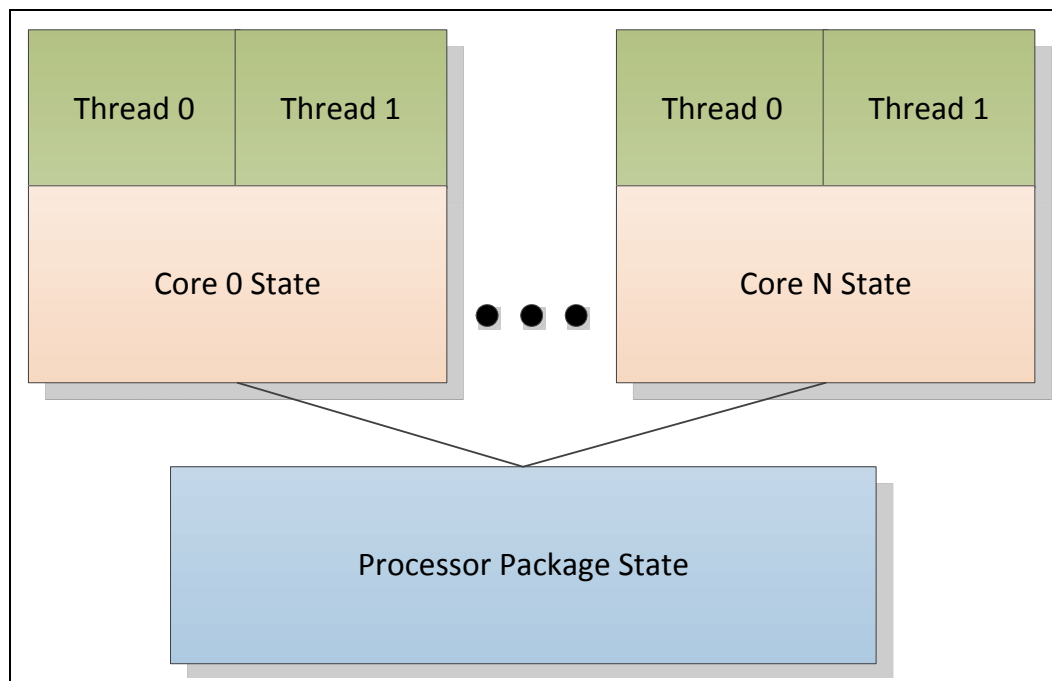
4.2.2 Low-Power Idle States

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, higher C-states have longer exit and entry latencies. Resolution of C-states occur at the thread, processor 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 4-2. Idle Power Management Breakdown of the Processor Cores



Entry and exit of the C-states at the thread and core level are shown in the following figure.

Figure 4-3. Thread and Core C-State Entry and Exit



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

Table 4-8. Coordination of Thread Power States at the Core Level

Processor Core C-State		Thread 1			
		C0	C1	C3	C6
Thread 0	C0	C0	C0	C0	C0
	C1	C0	C1 ¹	C1 ¹	C1 ¹
	C3	C0	C1 ¹	C3	C3
	C6	C0	C1 ¹	C3	C6

Note: If enabled, the core C-state will be C1E if all cores have resolved a core C1 state or higher.

4.2.3 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, must be enabled in the BIOS.

The BIOS can write to the C-state range field of the PMG_IO_CAPTURE MSR 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.

Note: 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.

4.2.4 Core C-State Rules

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

- A core C-state is determined by the lowest numerical thread state (such as Thread 0 requests C1E state while Thread 1 requests C3 state, resulting in a core C1E state). See the *G, S, and C Interface State Combinations* table.
- A 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 only that thread.
- If any thread in a core is in active (in C0 state), the core's C-state will resolve to C0 state.
- Any interrupt coming into the processor package may wake any core.



- A system reset re-initializes all processor cores.

Core C0 State

The normal operating state of a core where code is being executed.

Core C1/C1E State

C1/C1E is a low power state entered when all threads within a core execute a HLT or MWAIT(C1/C1E) instruction.

A System Management Interrupt (SMI) handler returns execution to either Normal state or the C1/C1E state. See the *Intel® 64 and IA-32 Architectures Software Developer's Manual* for more information.

While a core is in C1/C1E state, it processes bus snoops and snoops from other threads. For more information on C1E state, see [Section 4.2.5, "Package C-States."](#)

Core C3 State

Individual threads of a core can enter the C3 state by initiating a P_LVL2 I/O read to the P_BLK or an MWAIT(C3) instruction. A core in C3 state flushes the contents of its L1 instruction cache, L1 data cache, and L2 cache to the shared L3 cache, while maintaining its architectural state. All core clocks are stopped at this point. Because the core's caches are flushed, the processor does not wake any core that is in the C3 state when either a snoop is detected or when another core accesses cacheable memory.

Core C6 State

Individual threads of a core can enter the C6 state by initiating a P_LVL3 I/O read or an MWAIT(C6) instruction. Before entering core C6 state, the core will save its architectural state to a dedicated SRAM. Once complete, a core will have its voltage reduced to zero volts. During exit, the core is powered on and its architectural state is restored.

C-State Auto-Demotion

In general, deeper C-states, such as C6 state, 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 idle power. To increase residency and improve idle power in deeper C-states, the processor supports C-state auto-demotion.

There are two C-state auto-demotion options:

- C6 to C3 state
- C6/C3 To C1 state

The decision to demote a core from C6 to C3 or C3/C6 to C1 state is based on each core's immediate residency history and interrupt rate. If the interrupt rate experienced on a core is high and the residence in a deep C-state between such interrupts is low, the core can be demoted to a C3 or C1 state. A higher interrupt pattern is required to demote a core to C1 state as compared to C3 state.



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

4.2.5 Package C-States

The processor supports C0, C1/C1E, C3, C6 power 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 core C-state amongst all cores.
- A package C-state is automatically resolved by the processor depending on the core idle power states and the status of the platform components.
 - Each 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 state 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 core break event is received, the target core is activated and the break event message is forwarded to the target core.
 - If the break event is not masked, the target core enters the core C0 state and the processor enters package C0 state.
 - 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.

The following table shows package C-state resolution for a dual-core processor. The following figure summarizes package C-state transitions.




Table 4-9. Coordination of Core Power States at the Package Level

Package C-State		Core 1			
		C0	C1	C3	C6
Core 0	C0	C0	C0	C0	C0
	C1	C0	C1 ¹	C1 ¹	C1 ¹
	C3	C0	C1 ¹	C3	C3
	C6	C0	C1 ¹	C3	C6

Note:
 1. If enabled, the package C-state will be C1E if all cores have resolved a core C1 state or higher.

Figure 4-4. Package C-State Entry and Exit



Package C0 State

This is the normal operating state for the processor. The processor remains in the normal state when at least one of its cores is in the C0 or C1 state or when the platform has not granted permission to the processor to go into a low power state. Individual cores may be in lower power idle states while the package is in C0 state.



Package C1/C1E State

No additional power reduction actions are taken in the package C1 state. However, if the C1E sub-state is enabled, the processor automatically transitions to the lowest supported core clock frequency, followed by a reduction in voltage.

The package enters the C1 state low power state when:

- At least one core is in the C1 state.
- The other cores are in a C1 or deeper power state.

The package enters the C1E state when:

- All cores have directly requested C1E using MWAIT(C1) with a C1E sub-state hint.
- All cores are in a power state deeper than C1/C1E state but the package low power state is limited to C1/C1E using the PMG_CST_CONFIG_CONTROL MSR.
- All cores have requested C1 state using HLT or MWAIT(C1) and C1E auto-promotion is enabled in IA32_MISC_ENABLES.

No notification to the system occurs upon entry to C1/C1E state.

Package C2 State

Package C2 state is an internal processor state that cannot be explicitly requested by software. A processor enters Package C2 state when:

- All cores and graphics have requested a C3 or deeper power state, but constraints (LTR, programmed timer events in the near future, and so on) prevent entry to any state deeper than C 2 state. Or,
- All cores and graphics are in the C3 or deeper power states, and a memory access request is received. Upon completion of all outstanding memory requests, the processor transitions back into a deeper package C-state.

Package C3 State

A processor enters the package C3 low power state when:

- At least one core is in the C3 state.
- The other cores are in a C3 state or deeper power state and the processor has been granted permission by the platform.
- The platform has not granted a request to a package C6 or deeper state but has allowed a package C6 state.

In package C3 state, the L3 shared cache is valid.

Package C6 State

A processor enters the package C6 low-power state when:

- At least one core is in the C6 state.
- The other cores are in a C6 or deeper power state and the processor has been granted permission by the platform.
- If the cores are requesting C7 state, but the platform is limiting to a package C6 state, the last level cache in this case can be flushed.



In package C6 state all cores have saved their architectural state and have had their core voltages reduced to zero volts. It is possible the L3 shared cache is flushed and turned off in package C6 state. If at least one core is requesting C6 state, the L3 cache will not be flushed.

4.2.6 Package C-States and Display Resolution

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 4-10. Deepest Package C-State Available (Sheet 1 of 2)

Number of Displays ¹	Native Resolution	Deepest Available Package C-State
Single	800x600 60 Hz	PC6
Single	1024x768 60 Hz	PC6
Single	1280x1024 60 Hz	PC6
Single	1920x1080 60 Hz	PC6
Single	1920x1200 60 Hz	PC6
Single	1920x1440 60 Hz	PC6
Single	2048x1536 60 Hz	PC6
Single	2560x1600 60 Hz	PC6
Single	2560x1920 60 Hz	PC3
Single	2880x1620 60 Hz	PC3
Single	2880x1800 60 Hz	PC3
Single	3200x1800 60 Hz	PC3
Single	3200x2000 60 Hz	PC3
Single	3840x2160 60 Hz	PC3
Single	3840x2160 30 Hz	PC3
Single	4096x2160 24Hz	PC3
Multiple	800x600 60 Hz	PC6



Table 4-10. Deepest Package C-State Available (Sheet 2 of 2)

Number of Displays ¹	Native Resolution	Deepest Available Package C-State
Multiple	1024x768 60 Hz	PC6
Multiple	1280x1024 60 Hz	PC6
Multiple	1920x1080 60 Hz	PC3
Multiple	1920x1200 60 Hz	PC3
Multiple	1920x1440 60 Hz	PC3
Multiple	2048x1536 60 Hz	PC3
Multiple	2560x1600 60 Hz	PC2
Multiple	2560x1920 60 Hz	PC2
Multiple	2880x1620 60 Hz	PC2
Multiple	2880x1800 60 Hz	PC2
Multiple	3200x1800 60 Hz	PC2
Multiple	3200x2000 60 Hz	PC2
Multiple	3840x2160 60 Hz	PC2
Multiple	3840x2160 30 Hz	PC2
Multiple	4096x2160 24 Hz	PC2

Notes:
1. For multiple display cases, the resolution listed is the highest native resolution of all enabled displays, and PSR is internally disabled; that is, dual display with one 800x600 60 Hz display and one 2560x1600 60 Hz display will result in a deepest available package C-state of PC2.

4.3 Integrated Memory Controller (IMC) Power Management

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

4.3.1 Disabling Unused System Memory Outputs

Any system memory (SM) interface signal that goes to a memory module connector in which it is not connected to any actual memory devices (such as SO-DIMM 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 chip select and CKE signals are not driven.

At reset, all rows must 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 an SO-DIMM present, the SO-DIMM is not ensured to maintain data integrity.

CKE tri-state should be enabled by BIOS where appropriate, since at reset all rows must be assumed to be populated.



4.3.2 DRAM Power Management and Initialization

The processor implements extensive support for power management on the SDRAM interface. There are four SDRAM operations associated with the Clock Enable (CKE) signals, which the SDRAM controller supports. The processor drives four CKE pins to perform these operations.

The CKE is one of the power-save 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 three different types of power-down modes in package C0. The different power-down modes can be enabled through configuring "PM_PDWN_config_0_0_0_MCHBAR". 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:

1. **No power-down** (CKE disable)
2. **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 defined by tXP – small number of cycles. For this mode, DRAM DLL must be on.
3. **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 IDD2P1. 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 must be off.

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 the 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-offs 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_0_0_0_MCHBAR" is 6080h – that is, PPD/DLL-off mode with idle timer of 80h, or 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 DCKLs that a rank is idle that causes entry to the selected powermode. As this timer is set to a shorter time, the IMC will have more opportunities to put 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.

4.3.2.1 Initialization Role of CKE

During power-up, CKE is the only input to the SDRAM that has its level recognized (other than the DDR3L reset pin) once power is applied. It must 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.

4.3.2.2 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 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 core flushes pending cycles and then enters SDRAM ranks that are not used by Intel graphics memory 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.

4.3.2.3 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 core controller can be configured to put the devices in active power-down (CKE de-assertion with open pages) or pre-charge power-down (CKE de-assertion with all pages closed). Pre-charge 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.

4.3.2.4 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, ODE, 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 must be gated to prevent spurious results due to noise on the unused signals (typically handled automatically when input receiver is disabled).

4.3.3 DRAM Running Average Power Limitation (RAPL)

RAPL is a power and time constant pair. DRAM RAPL defines an average power constraint for the DRAM domain. Constraint is controlled by the PCU. Platform entities (PECI or in-band power driver) can specify a power limit for the DRAM domain. PCU continuously monitors the extent of DRAM throttling due to the power limit and re-budgets the limit between DIMMs.

4.3.4 DDR Electrical Power Gating (EPG)

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 V_{DDQ} for the majority of the logic to reduce idle power while keeping all critical DDR pins such as SM_DRAMRST#, CKE and VREF in the appropriate state.

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

4.4 PCI Express* Power Management

- Active power management is supported using L0s, and L1 states.
- All inputs and outputs disabled in L2/L3 Ready state.

4.5 Direct Media Interface (DMI) Power Management

Active power management is supported using L0s/L1 state.

4.6 Graphics Power Management

4.6.1 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 lower 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.

4.6.2 Graphics Render C-State

Render C-state (RC6) is a technique designed to optimize the average power to the graphics render engine during times of idleness. RC6 is entered when the graphics render engine, blitter engine, and the video engine have no workload being currently



worked on and no outstanding graphics memory transactions. When the idleness condition is met, the processor graphics will program the graphics render engine internal power rail into a low voltage state.

4.6.3 Intel® Graphics Dynamic Frequency

Intel Graphics Dynamic Frequency Technology is the ability of the processor and graphics cores to opportunistically increase frequency and/or voltage above the guaranteed processor and graphics frequency for the given part. Intel Graphics Dynamic Frequency Technology 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 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 try to place the graphics engine in the most energy efficient P-state.

§



5 Thermal Specifications Map

This chapter provides both component-level and system-level thermal management. Topics covered include processor thermal specifications, thermal profiles, thermal metrology, fan speed control, adaptive thermal monitor, THERMTRIP# signal, Digital Thermal Sensor (DTS), Intel Turbo Boost Technology, package power control, power plane control, and turbo time parameter.

The processor requires a thermal solution to maintain temperatures within its operating limits. Any attempt to operate the processor outside these operating limits may result in permanent damage to the processor and potentially other components within the system. Maintaining the proper thermal environment is key to reliable, long-term system operation.

A complete solution includes both component and system level thermal management features. Component level thermal solutions can include active or passive heatsinks attached to the processor integrated heat spreader (IHS).

To allow the optimal operation and long-term reliability of Intel processor-based systems, the processor must remain within the minimum and maximum case temperature (T_{CASE}) specifications as defined by the applicable thermal profile. Thermal solutions not designed to provide this level of thermal capability may affect the long-term reliability of the processor and system.

The processors implement a methodology for managing processor temperatures that is intended to support acoustic noise reduction through fan speed control and to assure processor reliability. Selection of the appropriate fan speed is based on the relative temperature data reported by the processor's Digital Temperature Sensor (DTS). The DTS can be read using the Platform Environment Control Interface (PECI) as described in Processor Temperature. Alternatively, when PEFI is monitored by the PCH, the processor temperature can be read from the PCH using the SMBus protocol defined in Embedded Controller Support Provided by the PCH. The temperature reported over PEFI is always a negative value and represents a delta below the onset of thermal control circuit (TCC) activation, as indicated by PROCHOT# (see Processor Temperature). Systems that implement fan speed control must be designed to use this data. Systems that do not alter the fan speed only need to ensure the case temperature meets the thermal profile specifications.

Analysis indicates that real applications are unlikely to cause the processor to consume maximum power dissipation for sustained time periods. Intel recommends that complete thermal solution designs target the Thermal Design Power (TDP), instead of the maximum processor power consumption. The Adaptive Thermal Monitor feature is intended to help protect the processor in the event that an application exceeds the TDP recommendation for a sustained time period. For more details on this feature, see Adaptive Thermal Monitor. To ensure maximum flexibility for future processors, systems should be designed to the Thermal Solution Capability guidelines, even if a processor with lower power dissipation is currently planned.



Table 5-1. Processor Thermal Specifications

Product	Max Power Package C1E (W) 1, 2, 5, 8	Max Power Package C3 (W) ^{1, 3, 5, 8}	Min Power Package C3 (W) ⁸	Max Power Package C6 (W) ^{1, 4, 5, 8}	Min Power Package C6 (W) ⁸	TTV Thermal Design Power (W) ^{6, 7, 10}	Min T _{CASE} (°C)	Max TTV T _{CASE} (°C)
Intel® Xeon® processor E3-1285 v4 80W 4C/8T, GT3	26	20	1.0	3.5	0	80	5	See Processor Thermal Profile Section
Intel® Xeon® processor E3-1285L v4 65W 4C/8T, GT3	23	17	1.0	3.5	0	65	5	See Processor Thermal Profile Section
Intel® Xeon® processor E3-1265L v4 35W 4C/8T, GT3	16	16	1.0	3.5	0	35	5	See Processor Thermal Profile Section

Notes:

- The package C-state power is the worst case power in the system configured as follows:
 - Memory configured for DDR3 1600 and populated with two DIMMs per channel.
 - DMI and PCIe links are at L1.
- Specification at DTS = 50 °C and minimum voltage loadline.
- Specification at DTS = 50 °C and minimum voltage loadline.
- Specification at DTS = 35 °C and minimum voltage loadline.
- These DTS values in Notes 2 – 4 are based on the TCC Activation MSR having a value of 100, see Processor Temperature.
- These values are specified at V_{CC_MAX} and V_{NOM} for all other voltage rails for all processor frequencies. Systems must be designed to ensure the processor is not to be subjected to any static V_{CC} and I_{CC} combination wherein V_{CCP} exceeds V_{CCP_MAX} at specified I_{CCP}. See the loadline specifications.
- Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at DTS = -1. TDP is achieved with the Memory configured for DDR3 1600 and 2 DIMMs per channel.
- Not 100% tested. Specified by design characterization.


5.1 Processor Thermal Profiles

This section provides thermal profiles for the processor families.



5.1.1 Processor (80W) Thermal Profile

Figure 5-1. Thermal Test Vehicle Thermal Profile for Processor (80W)



See the following table for discrete points that constitute the thermal profile.

Table 5-2. Thermal Test Vehicle Thermal Profile for Processor (80W) (Sheet 1 of 2)

Power (W)	T _{CASE_MAX} (°C)
Y = 0.33 * Power + 45	
0	45.00
2	45.66
4	46.32
6	46.98
8	47.64
10	48.30
12	48.96
14	49.62
16	50.28
18	50.94
20	51.60
22	52.26
24	52.92
26	53.58
28	54.24
30	54.90
32	55.56




Table 5-2. Thermal Test Vehicle Thermal Profile for Processor (80W) (Sheet 2 of 2)

Power (W)	T_{CASE_MAX} (°C)
34	56.22
36	56.88
38	57.54
40	58.20
42	58.86
44	59.52
46	60.18
48	60.84
50	61.50
52	62.16
54	62.82
56	63.48
58	64.14
60	64.80
62	65.46
64	66.12
66	66.78
68	67.44
70	68.10
72	68.76
74	69.42
76	70.08
78	70.74
80	71.40



5.1.2 Processor (65W) Thermal Profile

Figure 5-2. Thermal Test Vehicle Thermal Profile for Processor (65W)



See the following table for discrete points that constitute the thermal profile.

Table 5-3. Thermal Test Vehicle Thermal Profile for Processor (65W) (Sheet 1 of 2)

Power (W)	T _{CASE_MAX} (°C)
Y = 0.41 * Power + 44.7	
0	44.7
2	45.52
4	46.34
6	47.16
8	47.98
10	48.80
12	49.62
14	50.44
16	51.26
18	52.08
20	52.90
22	53.72
24	54.54
26	55.36
28	56.18
30	57.00
32	57.82




Table 5-3. Thermal Test Vehicle Thermal Profile for Processor (65W) (Sheet 2 of 2)

Power (W)	T _{CASE_MAX} (°C)
34	58.64
36	59.46
38	60.28
40	61.10
42	61.92
44	62.74
46	63.56
48	64.38
50	65.20
52	66.02
54	66.84
56	67.66
58	68.48
60	69.30
62	70.12
64	70.94
65	71.35

5.1.3 Processor (35W) Thermal Profile

Figure 5-3. Thermal Test Vehicle Thermal Profile for Processor (35W)



See the following table for discrete points that constitute the thermal profile.





Table 5-4. Thermal Test Vehicle Thermal Profile for Processor (35W)

Power (W)	T _{CASE_MAX} (°C)
Y = 0.51 * Power + 48.5	
0	48.50
2	49.52
4	50.54
6	51.56
8	52.58
10	53.60
12	54.62
14	55.64
16	56.66
18	57.68
20	58.70
22	59.72
24	60.74
26	61.76
28	62.78
30	63.80
32	64.82
34	65.84
35	66.86

5.2 Thermal Metrology

The maximum Thermal Test Vehicle (TTV) case temperatures (T_{CASE-MAX}) can be derived from the data in the appropriate TTV thermal profile earlier in this chapter. The TTV T_{CASE} is measured at the geometric top center of the TTV integrated heat spreader (IHS). The following figure illustrates the location where T_{CASE} temperature measurements should be made.

Figure 5-4. Thermal Test Vehicle (TTV) Case Temperature (T_{CASE}) Measurement Location



Note: THERM-X OF CALIFORNIA can machine the groove and attach a thermocouple to the IHS. The supplier is subject to change without notice. THERM-X OF CALIFORNIA, 1837 Whipple Road, Hayward, Ca 94544. Skip Johnson +1-510-441-7566 Ext. 242 skipj@therm-x.com. The vendor part number is XTMS1565.

5.3 Fan Speed Control Scheme with Digital Thermal Sensor (DTS) 1.1

To correctly use DTS 1.1, the designer must first select a worst case scenario $T_{AMBIENT}$ and ensure that the Fan Speed Control (FSC) can provide a Ψ_{CA} that is equivalent or greater than the Ψ_{CA} specification.

The DTS 1.1 implementation consists of two points: a Ψ_{CA} at $T_{CONTROL}$ and a Ψ_{CA} at $DTS = -1$.

The Ψ_{CA} point at $DTS = -1$ defines the minimum Ψ_{CA} required at TDP considering the worst case system design $T_{AMBIENT}$ design point:

$$\Psi_{CA} = (T_{CASE-MAX} - T_{AMBIENT-TARGET}) / TDP$$

For example, for a 95 W TDP part, the T_{case} maximum is 72.6 °C and at a worst case design point of 40 °C local ambient this will result in:



$$\Psi_{CA} = (72.6 - 40) / 95 = 0.34 \text{ }^{\circ}\text{C/W}$$

Similarly for a system with a design target of 45 °C ambient, the Ψ_{CA} at DTS = -1 needed will be 0.29 °C/W.

The second point defines the thermal solution performance (Ψ_{CA}) at $T_{CONTROL}$. The following table lists the required Ψ_{CA} for the various TDP processors.

These two points define the operational limits for the processor for DTS 1.1 implementation. At $T_{CONTROL}$ the fan speed must be programmed such that the resulting Ψ_{CA} is better than or equivalent to the required Ψ_{CA} listed in the following table. Similarly, the fan speed should be set at DTS = -1 such that the thermal solution performance is better than or equivalent to the Ψ_{CA} requirements at $T_{AMBIENT-MAX}$. The fan speed controller must linearly ramp the fan speed from processor DTS = $T_{CONTROL}$ to processor DTS = -1.

Figure 5-5. Digital Thermal Sensor (DTS) 1.1 Definition Points




Table 5-5. Digital Thermal Sensor (DTS) 1.1 Thermal Solution Performance Above $T_{CONTROL}$ (Sheet 1 of 2)

Processor TDP	Ψ_{CA} at DTS = $T_{CONTROL}$ ^{1, 2} At System $T_{AMBIENT-MAX}$ = 30 °C	Ψ_{CA} at DTS = -1 At System $T_{AMBIENT-MAX}$ = 40 °C	Ψ_{CA} at DTS = -1 At System $T_{AMBIENT-MAX}$ = 45 °C	Ψ_{CA} at DTS = -1 At System $T_{AMBIENT-MAX}$ = 50 °C
80 W	0.619	0.387	0.330	0.273



Table 5-5. Digital Thermal Sensor (DTS) 1.1 Thermal Solution Performance Above T_{CONTROL} (Sheet 2 of 2)

Processor TDP	Ψ_{CA} at DTS = T _{CONTROL} ^{1, 2} At System T _{AMBIENT-MAX} = 30 °C	Ψ_{CA} at DTS = -1 At System T _{AMBIENT-MAX} = 40 °C	Ψ_{CA} at DTS = -1 At System T _{AMBIENT-MAX} = 45 °C	Ψ_{CA} at DTS = -1 At System T _{AMBIENT-MAX} = 50 °C
65 W	0.684	0.481	0.404	0.327
35 W	1.406	0.753	0.610	0.467

Notes:

- Ψ_{CA} at "DTS = T_{CONTROL}" is applicable to systems that have an internal T_{RISE} (T_{ROOM} temperature to Processor cooling fan inlet) of less than 10 °C. In case the expected T_{RISE} is greater than 10 °C, a correction factor should be used as explained below. For each 1 °C T_{RISE} above 10 °C, the correction factor (CF) is defined as CF = 1.7 / (processor TDP)
- Example: A chassis T_{RISE} assumption is 12 °C for a 95 W TDP processor:
 CF = 1.7 / 95 W = 0.018 /W
 For T_{RISE} > 10 °C
 Ψ_{CA} at T_{CONTROL} = (Value provide in Column 2) - (T_{RISE} - 10) * CF
 Ψ_{CA} = 0.627 - (12 - 10) * 0.018 = 0.591 °C/W
 In this case, the fan speed should be set slightly higher, equivalent to Ψ_{CA} = 0.591 °C/W

5.4 Fan Speed Control Scheme with Digital Thermal Sensor (DTS) 2.0

To simplify processor thermal specification compliance, the processor calculates the DTS Thermal Profile from T_{CONTROL} Offset, TCC Activation Temperature, TDP, and the Thermal Margin Slope provided in the following table.

Note: TCC Activation Offset is 0 for the processors.

Using the DTS Thermal Profile, the processor can calculate and report the Thermal Margin, where a value less than 0 indicates that the processor needs additional cooling, and a value greater than 0 indicates that the processor is sufficiently cooled. Refer to the processor Thermal Mechanical Design Guidelines (TMDG) for additional information (see Related Documents section).



Figure 5-6. Digital Thermal Sensor (DTS) Thermal Profile Definition




Table 5-6. Thermal Margin Slope

Die Configuration	TDP (W)	TCC Activation Temperature (°C) MSR 1A2h 23:16	Temperature Control Offset MSR 1A2h 15:8	Thermal Margin Slope (°C / W)
4+3e	80	100	20	0.671
4+3e	65	92	6	0.722
4+3e	35	75	6	0.806

5.5 Product Family Thermal Specifications

This section provides thermal specifications (Thermal Profile) and design guidelines for enabled thermal solutions to cool the processor.

5.5.1 Performance Targets

The following table provides boundary conditions and performance targets as guidance for thermal solution design. Thermal solutions must be able to comply with the Maximum T_{CASE} Thermal Profile.



Table 5-7. Boundary Conditions, Performance Targets, and T_{CASE} Specifications

Processor	Package TDP	Heatsink ¹	T _{LA} , Airflow, RPM, Ψ _{CA} ²	Maximum T _{CASE} Thermal Profile ³
4C/GT3 80 W	80 W	1U AI Heatpipe	40 °C, 11.5 CFM, 0.383 °C/W	y = 0.33 * Power + 45.0
4C/GT3 65 W	65 W	1U AI Heatpipe	40 °C, 11.5 CFM, 0.487 °C/W	y = 0.41 * Power + 44.7
4C/GT3 35 W	35 W	1U AI Heatpipe	45.4 °C, 11.5 CFM, 0.597 °C/W	y = 0.51 * Power + 48.5

1. Thermal Solution information can be found in the following table. Alternate heatsink for 65W is Active AI Core (DHA-B), up to 40C. Alternate heatsink for 35W is Active Short (DHA-D), up to 45.4C.

2. These boundary conditions and performance targets are used to generate processor thermal specifications and to provide guidance for heatsink design. Values are for the heatsink shown in the adjacent column are calculated at sea level, and are expected to meet the Thermal Profile at TDP. T_{LA} is the local ambient temperature of the heatsink inlet air. Airflow is through the heatsink fins with zero bypass for a passive heatsink. RPM is fan revolutions per minute for an active heatsink. Ψ_{CA} is the maximum target (mean + 3 sigma) for the thermal characterization parameter. For more information on the thermal characterization parameter, refer to the processor Thermal Mechanical Design Guidelines (see Related Documents section).

3. Maximum T_{CASE} Thermal Profile is the specification that must be complied to. Any Attempt to operate the processor outside these operating limits may result in permanent damage to the processor and potentially other system components.

5.6 Processor Temperature

A software readable field in the TEMPERATURE_TARGET register that contains the minimum temperature at which the TCC will be activated and PROCHOT# will be asserted. The TCC activation temperature is calibrated on a part-by-part basis and normal factory variation may result in the actual TCC activation temperature being higher than the value listed in the register. TCC activation temperatures may change based on processor stepping, frequency or manufacturing efficiencies.

5.7 Adaptive Thermal Monitor

The Adaptive Thermal Monitor feature provides an enhanced method for controlling the processor temperature when the processor silicon exceeds the Thermal Control Circuit (TCC) activation temperature. Adaptive Thermal Monitor uses TCC activation to reduce processor power using a combination of methods. The first method (Frequency control, similar to Thermal Monitor 2 (TM2) in previous generation processors) involves the processor reducing its operating frequency (using the core ratio multiplier) and internal core voltage. This combination of lower frequency and core voltage results in a reduction of the processor power consumption. The second method (clock modulation, known as Thermal Monitor 1 or TM1 in previous generation processors) reduces power consumption by modulating (starting and stopping) the internal processor core clocks. The processor intelligently selects the appropriate TCC method to use on a dynamic basis. BIOS is not required to select a specific method (as with previous-generation processors supporting TM1 or TM2). The temperature at which Adaptive Thermal Monitor activates the Thermal Control Circuit is factory calibrated and is not user configurable. Snooping and interrupt processing are performed in the normal manner while the TCC is active.



When the TCC activation temperature is reached, the processor will initiate TM2 in attempt to reduce its temperature. If TM2 is unable to reduce the processor temperature, TM1 will be also be activated. TM1 and TM2 will work together (clocks will be modulated at the lowest frequency ratio) to reduce power dissipation and temperature.

With a properly designed and characterized thermal solution, it is anticipated that the TCC will only be activated 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. An under-designed thermal solution that is not able to prevent excessive activation of the TCC in the anticipated ambient environment may cause a noticeable performance loss, and in some cases may result in a T_{CASE} that exceeds the specified maximum temperature and may affect the long-term reliability of the processor. In addition, a thermal solution that is significantly under designed may not be capable of cooling the processor even when the TCC is active continuously. See the appropriate processor Thermal Mechanical Design Guidelines for information on designing a compliant thermal solution.

The Thermal Monitor does not require any additional hardware, software drivers, or interrupt handling routines. The following sections provide more details on the different TCC mechanisms used by the processor.

5.7.1 Frequency Control

When the Digital Temperature Sensor (DTS) reaches a value of 0 (DTS temperatures reported using PECCI may not equal zero when PROCHOT# is activated), the TCC will be activated and the PROCHOT# signal will be asserted if configured as bi-directional. This indicates the processor temperature has met or exceeded the factory calibrated trip temperature and it will take action to reduce the temperature.

Upon activation of the TCC, the processor will stop the core clocks, reduce the core ratio multiplier by 1 ratio and restart the clocks. All processor activity stops during this frequency transition that occurs within 2 μ s. Once the clocks have been restarted at the new lower frequency, processor activity resumes while the core voltage is reduced by the internal voltage regulator. Running the processor at the lower frequency and voltage will reduce power consumption and should allow the processor to cool off. If after 1 ms the processor is still too hot (the temperature has not dropped below the TCC activation point, DTS still = 0 and PROCHOT is still active), then a second frequency and voltage transition will take place. This sequence of temperature checking and frequency and voltage reduction will continue until either the minimum frequency has been reached or the processor temperature has dropped below the TCC activation point.

If the processor temperature remains above the TCC activation point even after the minimum frequency has been reached, then clock modulation (described below) at that minimum frequency will be initiated.

There is no end user software or hardware mechanism to initiate this automated TCC activation behavior.

A small amount of hysteresis has been included to prevent rapid active/inactive transitions of the TCC when the processor temperature is near the TCC activation temperature. Once the temperature has dropped below the trip temperature and the hysteresis timer has expired, the operating frequency and voltage transition back to



the normal system operating point using the intermediate VID/frequency points. Transition of the VID code will occur first, to insure proper operation as the frequency is increased.

5.7.2 Clock Modulation

Clock modulation is a second method of thermal control available to the processor. Clock modulation is performed by rapidly turning the clocks off and on at a duty cycle that should reduce power dissipation by about 50% (typically a 30–50% duty cycle). Clocks often will not be off for more than 32 microseconds when the TCC is active. Cycle times are independent of processor frequency. The duty cycle for the TCC, when activated by the Thermal Monitor, is factory configured and cannot be modified.

A small amount of hysteresis has been included to prevent rapid active/inactive transitions of the TCC 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 TCC goes inactive and clock modulation ceases.

5.7.3 Immediate Transition to Combined TM1 and TM2

When the TCC is activated, the processor will sequentially step down the ratio multipliers and VIDs in an attempt to reduce the silicon temperature. If the temperature continues to increase and exceeds the TCC activation temperature by approximately 5 °C before the lowest ratio/VID combination has been reached, the processor will immediately transition to the combined TM1/TM2 condition. The processor remains in this state until the temperature has dropped below the TCC activation point. Once below the TCC activation temperature, TM1 will be discontinued and TM2 will be exited by stepping up to the appropriate ratio/VID state.

5.7.4 Critical Temperature Flag

If TM2 is unable to reduce the processor temperature, then TM1 will be also be activated. TM1 and TM2 will then work together to reduce power dissipation and temperature. It is expected that only a catastrophic thermal solution failure would create a situation where both TM1 and TM2 are active.

If TM1 and TM2 have both been active for greater than 20 ms and the processor temperature has not dropped below the TCC activation point, the Critical Temperature Flag in the IA32_THERM_STATUS MSR will be set. This flag is an indicator of a catastrophic thermal solution failure and that the processor cannot reduce its temperature. Unless immediate action is taken to resolve the failure, the processor will probably reach the Thermtrip temperature (see [Section 6.8](#), , “Testability Signals”) within a short time. To prevent possible permanent silicon damage, Intel recommends removing power from the processor within ½ second of the Critical Temperature Flag being set.

5.7.5 PROCHOT# Signal

An external signal, PROCHOT# (processor hot), is asserted when the processor core temperature has exceeded its specification. If Adaptive Thermal Monitor is enabled (it must be enabled for the processor to be operating within specification), the TCC will be active when PROCHOT# is asserted.



The processor can be configured to generate an interrupt upon the assertion or de-assertion of PROCHOT#.

By default, the PROCHOT# signal is set to bi-directional. However, it is recommended to configure the signal as an input only. When configured as an input or bi-directional signal, PROCHOT# can be used for thermally protecting other platform components should they overheat as well. When PROCHOT# is driven by an external device:

- The package will immediately transition to the minimum operation points (voltage and frequency) supported by the processor and graphics cores. This is contrary to the internally-generated Adaptive Thermal Monitor response.
- Clock modulation is not activated.

The TCC will remain active until the system de-asserts PROCHOT#. The processor can be configured to generate an interrupt upon assertion and de-assertion of the PROCHOT# signal.

Note: Toggling PROCHOT# more than once in 1.5 ms period will result in constant Pn state of the processor.

Note: A corner case exists for PROCHOT# configured as a bi-directional signal that can cause several milliseconds of delay to a system assertion of PROCHOT# when the output function is asserted.

As an output, PROCHOT# (Processor Hot) will go active when the processor temperature monitoring sensor detects that one or more cores has reached its maximum safe operating temperature. This indicates that the processor Thermal Control Circuit (TCC) has been activated, if enabled. As an input, assertion of PROCHOT# by the system will activate the TCC for all cores. TCC activation when PROCHOT# is asserted by the system will result in the processor immediately transitioning to the minimum frequency and corresponding voltage (using Frequency control). Clock modulation is not activated in this case. The TCC will remain active until the system de-asserts PROCHOT#.

Use of PROCHOT# in input or bi-directional mode can allow VR thermal designs to target maximum sustained current instead of maximum current. Systems should still provide proper cooling for the Voltage Regulator (VR), and rely on PROCHOT# only as a backup in case of system cooling failure. The system thermal design should allow the power delivery circuitry to operate within its temperature specification even while the processor is operating at its Thermal Design Power.

5.8 THERMTRIP# Signal

Regardless of whether or not Adaptive Thermal Monitor is enabled, in the event of a catastrophic cooling failure, the processor will automatically shut down when the silicon has reached an elevated temperature (refer to the THERMTRIP# definition in [Section 6.9](#), "Error and Thermal Protection Signals"). At this point, the THERMTRIP# signal will go active and stay active. THERMTRIP# activation is independent of processor activity. The temperature at which THERMTRIP# asserts is not user configurable and is not software visible.



5.9 Digital Thermal Sensor

Each processor execution core has an on-die Digital Thermal Sensor (DTS) that detects the core's instantaneous temperature. The DTS is the preferred method of monitoring processor die temperature because:

- It is located near the hottest portions of the die.
- It can accurately track the die temperature and ensure that the Adaptive Thermal Monitor is not excessively activated.

Temperature values from the DTS can be retrieved through:

- A software interface using processor Model Specific Register (MSR).
- A processor hardware interface.

When temperature is retrieved by the processor MSR, it is the instantaneous temperature of the given core. 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 within the PACKAGE_THERM_STATUS MSR 1B1h and IA32_THERM_STATUS MSR 19Ch.

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_{jMAX}), 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 in the TEMPERATURE_TARGET MSR 1A2h. The temperature returned by the DTS is an implied negative integer indicating the relative offset from T_{jMAX} . The DTS does not report temperatures greater than T_{jMAX} . 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 0h, 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 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 MSRs. These thresholds have the capability of generating interrupts using the core's local APIC. Refer to the *Intel® 64 and IA-32 Architectures Software Developer's Manual* for specific register and programming details.

5.9.1 Digital Thermal Sensor Accuracy (Taccuracy)

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



5.10 Intel® Turbo Boost Technology Thermal Considerations

Intel Turbo Boost Technology allows processor cores and integrated graphics cores to run faster than the baseline frequency. During a turbo event, the processor can exceed its TDP power for brief periods. Turbo is invoked opportunistically and automatically as long as the processor is conforming to its temperature, power delivery, and current specification limits. Thus, thermal solutions and platform cooling that are designed to less than thermal design guidance may experience thermal and performance issues since more applications will tend to run at or near the maximum power limit for significant periods of time.

5.11 Intel® Turbo Boost Technology Power Control and Reporting

Package processor core and internal graphics core powers are self monitored and correspondingly reported out.

- With the processor turbo disabled, rolling average power over 5 seconds will not exceed the TDP rating of the part for typical applications.
- With turbo enabled (see [Figure 5-1, “Thermal Test Vehicle Thermal Profile for Processor \(80W\)”](#))
 - For the PL1: Package rolling average of the power set in POWER_LIMIT_1 (TURBO_POWER_LIMIT MSR 0610h bits [14:0]) over time window set in POWER_LIMIT_1_TIME (TURBO_POWER_LIMIT MSR 0610h bits [23:17]) must be less than or equal to the TDP package power as read from the PACKAGE_POWER_SKU MSR 0614h for typical applications. Power control is valid only when the processor is operating in turbo. PL1 lower than the package TDP is not guaranteed.
 - For the PL2: Package power will be controlled to a value set in POWER_LIMIT_2 (TURBO_POWER_LIMIT MSR 0610h bits [46:32]). Occasional brief power excursions may occur for periods of less than 10 ms over PL2.

The processor monitors its own power consumption to control turbo behavior, assuming the following:

- The power monitor is not 100% tested across all processors.
- The Power Limit 2 (PL2) control is only valid for power levels set at or above TDP and under workloads with similar activity ratios as the product TDP workload. This also assumes the processor is working within other product specifications.
- Setting power limits (PL1 or PL2) below TDP are not ensured to be followed, and are not characterized for accuracy.
- Under unknown work loads and unforeseen applications the average processor power may exceed Power Limit 1 (PL1).
- Uncharacterized workloads may exist that could result in higher turbo frequencies and power. If that were to happen, the processor Thermal Control Circuitry (TCC) would protect the processor. The TCC protection must be enabled by the platform for the product to be within specification.

An illustration of Intel Turbo Boost Technology power control is shown in the following sections and figures. Multiple controls operate simultaneously allowing for customization for multiple system thermal and power limitations. These controls provide turbo optimizations within system constraints.

5.11.1 Package Power Control

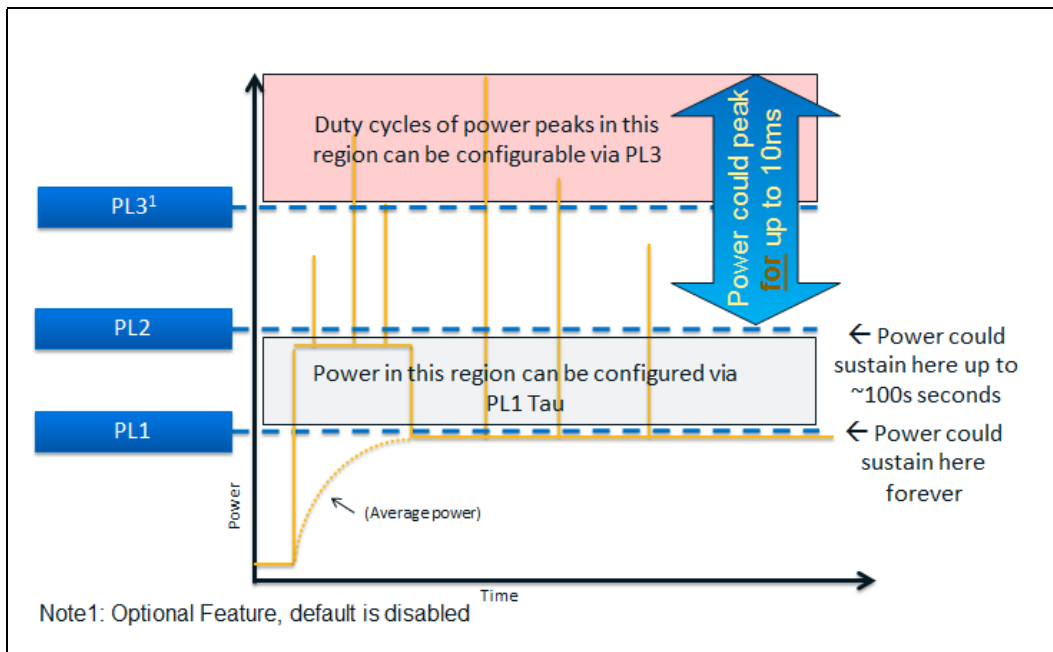
The package power control settings of PL1, PL2, and PL3 Tau allow the designer to configure Intel Turbo Boost Technology 2.0 to match the platform power delivery and package thermal solution limitations.

- Power Limit 1 (PL1): A threshold for average power that will not exceed; recommend to set equal to 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.
- Turbo Time Parameter (Tau): An averaging constant used for PL1 exponential weighted moving average (EWMA) power calculation.

Note: Implementation of Intel Turbo Boost Technology 2.0 only requires configuring PL1, PL1 Tau, and PL2.

Note: PL3 is disabled by default.

Figure 5-7. Package Power Control



§



6 Signal Description

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

Notation	Signal Type
I	Input pin
O	Output pin
I/O	Bi-directional Input/Output pin

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

Table 6-1. Signal Description Buffer Types

Signal	Description
PCI Express*	PCI Express* interface signals. These signals are compatible with PCI Express 3.0 Signaling Environment AC Specifications and are AC coupled. The buffers are not 3.3 V- tolerant. See the <i>PCI Express Base Specification 3.0</i> .
DMI	Direct Media Interface signals. These signals are compatible with PCI Express 2.0 Signaling Environment AC Specifications, but are DC coupled. The buffers are not 3.3 V- tolerant.
CMOS	CMOS buffers. 1.05 V- tolerant
DDR3L	DDR3L buffers: 1.35 V- 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
Asynchronous ¹	Signal has no timing relationship with any reference clock.
1. Qualifier for a buffer type.	

6.1 System Memory Interface Signals

Table 6-2. Memory Channel A Signals (Sheet 1 of 2)

Signal Name	Description	Direction / Buffer Type
SA_BS[2:0]	Bank Select: These signals define which banks are selected within each SDRAM rank.	O DDR3L
SA_WE#	Write Enable Control Signal: This signal is used with SA_RAS# and SA_CAS# (along with SA_CS#) to define the SDRAM Commands.	O DDR3L
SA_RAS#	RAS Control Signal: This signal is used with SA_CAS# and SA_WE# (along with SA_CS#) to define the SRAM Commands.	O DDR3L
SA_CAS#	CAS Control Signal: This signal is used with SA_RAS# and SA_WE# (along with SA_CS#) to define the SRAM Commands.	O DDR3L



Table 6-2. Memory Channel A Signals (Sheet 2 of 2)

Signal Name	Description	Direction / Buffer Type
SA_DQS[8:0] SA_DQSN[8:0]	Data Strobes: SA_DQS[8:0] and its complement signal group make up a differential strobe pair. The data is captured at the crossing point of SA_DQS[8:0] and SA_DQS#[8:0] during read and write transactions.	I/O DDR3L
SA_DQ[63:0]	Data Bus: Channel A data signal interface to the SDRAM data bus.	I/O DDR3L
SA_ECC_CB[7:0]	ECC Data Lines: Data Lines for ECC Check Byte.	I/O DDR3L
SA_MA[15:0]	Memory Address: These signals are used to provide the multiplexed row and column address to the SDRAM.	O DDR3L
SA_CK[3:0]	SDRAM Differential Clock: These signals are Channel A SDRAM Differential clock signal pairs. The crossing of the positive edge of SA_CK and the negative edge of its complement SA_CK# are used to sample the command and control signals on the SDRAM.	O DDR3L
SA_CKE[3:0]	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 DDR3L
SA_CS#[3:0]	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 DDR3L
SA_ODT[3:0]	On Die Termination: Active Termination Control.	O DDR3L

Table 6-3. Memory Channel B Signals (Sheet 1 of 2)

Signal Name	Description	Direction / Buffer Type
SB_BS[2:0]	Bank Select: These signals define which banks are selected within each SDRAM rank.	O DDR3L
SB_WE#	Write Enable Control Signal: This signal is used with SB_RAS# and SB_CAS# (along with SB_CS#) to define the SDRAM Commands.	O DDR3L
SB_RAS#	RAS Control Signal: This signal is used with SB_CAS# and SB_WE# (along with SB_CS#) to define the SRAM Commands.	O DDR3L
SB_CAS#	CAS Control Signal: This signal is used with SB_RAS# and SB_WE# (along with SB_CS#) to define the SRAM Commands.	O DDR3L
SB_DQS[8:0] SB_DQSN[8:0]	Data Strobes: SB_DQS[8:0] and its complement signal group make up a differential strobe pair. The data is captured at the crossing point of SB_DQS[8:0] and its SB_DQS#[8:0] during read and write transactions.	I/O DDR3L
SB_DQ[63:0]	Data Bus: Channel B data signal interface to the SDRAM data bus.	I/O DDR3L
SB_ECC_CB[7:0]	ECC Data Lines: Data Lines for ECC Check Byte.	I/O DDR3L
SB_MA[15:0]	Memory Address: These signals are used to provide the multiplexed row and column address to the SDRAM.	O DDR3L
SB_CK[3:0]	SDRAM Differential Clock: Channel B SDRAM Differential clock signal pair. The crossing of the positive edge of SB_CK and the negative edge of its complement SB_CK# are used to sample the command and control signals on the SDRAM.	O DDR3L



Table 6-3. Memory Channel B Signals (Sheet 2 of 2)

Signal Name	Description	Direction / Buffer Type
SB_CKE[3:0]	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 DDR3L
SB_CS#[3:0]	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 DDR3L
SB_ODT[3:0]	On Die Termination: Active Termination Control.	O DDR3L

6.2 Memory Reference and Compensation

Table 6-4. Memory Reference and Compensation

Signal Name	Description	Direction / Buffer Type
SM_RCOMP[2:0]	System Memory Impedance Compensation:	I A
SM_VREF	DDR3L Reference Voltage: This signal is used as a reference voltage to the DDR3L controller and is defined as $V_{DDQ}/2$	O DDR3L
SA_DIMM_VREFDQ SB_DIMM_VREFDQ	Memory Channel A/B DIMM DQ Voltage Reference: The output pins are connected to the DIMMs, and holds $V_{DDQ}/2$ as reference voltage.	O DDR3L



6.3 Reset and Miscellaneous Signals

Table 6-5. Reset and Miscellaneous Signals

Signal Name	Description	Direction / Buffer Type
CFG[19:0]	<p>Configuration Signals: The CFG signals have a default value of '1' if not terminated on the board.</p> <ul style="list-style-type: none"> • CFG[1:0]: Reserved configuration lane. A test point may be placed on the board for these lanes. • CFG[2]: PCI Express* Static x16 Lane Numbering Reversal. <ul style="list-style-type: none"> – 1 = Normal operation – 0 = Lane numbers reversed. • CFG[3]: MSR Privacy Bit Feature <ul style="list-style-type: none"> – 1 = Debug capability is determined by IA32_Debug_Interface_MSR (C80h) bit[0] setting – 0 = IA32_Debug_Interface_MSR (C80h) bit[0] default setting overridden • CFG[4]: Reserved configuration lane. A test point may be placed on the board for this lane. • CFG[6:5]: PCI Express* Bifurcation: ¹ <ul style="list-style-type: none"> – 00 = 1 x8, 2 x4 PCI Express* – 01 = reserved – 10 = 2 x8 PCI Express* – 11 = 1 x16 PCI Express* • CFG[19:7]: Reserved configuration lanes. A test point may be placed on the board for these lands. 	I/O GTL
CFG_RCOMP	Configuration resistance compensation. Use a 49.9 Ω ±1% resistor to ground.	–
FC_x	FC (Future Compatibility) signals are signals that are available for compatibility with other processors. A test point may be placed on the board for these lands.	
PM_SYNC	Power Management Sync: A sideband signal to communicate power management status from the platform to the processor.	I CMOS
PWR_DEBUG#	Signal is for debug. R	I Asynchronous CMOS
IST_TRIGGER	Signal is for IFDIM testing only.	I CMOS
IVR_ERROR	Signal is for debug. If both THERMTRIP# and this signal are simultaneously asserted, the processor has encountered an unrecoverable power delivery fault and has engaged automatic shutdown as a result.	O CMOS
RESET#	Platform Reset pin driven by the PCH.	I CMOS
RSVD RSVD_TP RSVD_NCTF	RESERVED: All signals that are RSVD and RSVD_NCTF must be left unconnected on the board. Intel recommends that all RSVD_TP signals have via test points.	No Connect Test Point Non-Critical to Function
SM_DRAMRST#	DRAM Reset: Reset signal from processor to DRAM devices. One signal common to all channels.	O CMOS
TESTLO_x	TESTLO should be individually connected to V _{SS} through a resistor.	
Note:		
1. PCIe bifurcation support varies with the processor and PCH SKUs used.		



6.4 PCI Express*-Based Interface Signals

Table 6-6. PCI Express* Graphics Interface Signals

Signal Name	Description	Direction / Buffer Type
PEG_RCOMP	PCI Express Resistance Compensation	I A
PEG_RXP[15:0] PEG_RXN[15:0]	PCI Express Receive Differential Pair	I PCI Express
PEG_TXP[15:0] PEG_TXN[15:0]	PCI Express Transmit Differential Pair	O PCI Express

6.5 Display Interface Signals

Table 6-7. Display Interface Signals

Signal Name	Description	Direction / Buffer Type
FDI_TXP[1:0] FDI_TXN[1:0]	Intel Flexible Display Interface Transmit Differential Pair	O FDI
DDIB_TXP[3:0] DDIB_TXN[3:0]	Digital Display Interface Transmit Differential Pair	O FDI
DDIC_TXP[3:0] DDIC_TXN[3:0]	Digital Display Interface Transmit Differential Pair	O FDI
DDID_TXP[3:0] DDID_TXN[3:0]	Digital Display Interface Transmit Differential Pair	O FDI
FDI_CSNC	Intel Flexible Display Interface Sync	I CMOS
DISP_INT	Intel Flexible Display Interface Hot-Plug Interrupt	I Asynchronous CMOS

6.6 Direct Media Interface (DMI)

Table 6-8. Direct Media Interface (DMI) – Processor to PCH Serial Interface

Signal Name	Description	Direction / Buffer Type
DMI_RXP[3:0] DMI_RXN[3:0]	DMI Input from PCH: Direct Media Interface receive differential pair.	I DMI
DMI_TXP[3:0] DMI_TXN[3:0]	DMI Output to PCH: Direct Media Interface transmit differential pair.	O DMI



6.7 Phase Locked Loop (PLL) Signals

Table 6-9. Phase Locked Loop (PLL) Signals

Signal Name	Description	Direction / Buffer Type
BCLKP BCLKN	Differential bus clock input to the processor	I Diff Clk
DPLL_REF_CLKP DPLL_REF_CLKN	Embedded Display Port PLL Differential Clock In: 135 MHz	I Diff Clk
SSC_DPLL_REF_CLKP SSC_ DPLL_REF_CLKN	Spread Spectrum Embedded DisplayPort PLL Differential Clock In: 135 MHz	I Diff Clk

6.8 Testability Signals

Table 6-10. Testability Signals

Signal Name	Description	Direction / Buffer Type
BPM#[7:0]	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 CMOS
DBR#	Debug Reset: This signal is used only in systems where no debug port is implemented on the system board. DBR# is used by a debug port interposer so that an in-target probe can drive system reset.	O
PRDY#	Processor Ready: This signal is a processor output used by debug tools to determine processor debug readiness.	O Asynchronous CMOS
PREQ#	Processor Request: This signal is used by debug tools to request debug operation of the processor.	I Asynchronous CMOS
TCK	Test Clock: This signal provides the clock input for the processor Test Bus (also known as the Test Access Port). This signal must be driven low or allowed to float during power on Reset.	I GTL
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
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 Open Drain
TMS	Test Mode Select: This is a JTAG specification supported signal used by debug tools.	I GTL
TRST#	Test Reset: This signal resets the Test Access Port (TAP) logic. This signal must be driven low during power on Reset.	I GTL



6.9 Error and Thermal Protection Signals

Table 6-11. Error and Thermal Protection Signals

Signal Name	Description	Direction / Buffer Type
CATERR#	Catastrophic Error: This signal indicates that the system has experienced a catastrophic error and cannot continue to operate. The processor will set this 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 GTL
PECI	Platform Environment Control Interface: A serial sideband interface to the processor, it is used primarily for thermal, power, and error management.	I/O Asynchronous
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.	GTL Input Open-Drain Output
THERMTRIP#	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 execution when the junction temperature exceeds approximately 130 °C. This is signaled to the system by the THERMTRIP# pin.	O Asynchronous OD Asynchronous CMOS

6.10 Power Sequencing Signals

Table 6-12. Power Sequencing Signals

Signal Name	Description	Direction / Buffer Type
SM_DRAMPWROK	SM_DRAMPWROK Processor Input: This signal connects to the PCH DRAMPWROK.	I Asynchronous CMOS
PWRGOOD	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 must then transition monotonically to a high state.	I Asynchronous CMOS
SKTOCC#	SKTOCC# (Socket Occupied)/PROC_DETECT# (Processor Detect): This signal is 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.	—



6.11 Processor Power Signals

Table 6-13. Processor Power Signals

Signal Name	Description	Direction / Buffer Type
VCC	Processor core power rail.	Ref
VCCIO_OUT	Processor power reference for I/O.	Ref
VCCST	Sustain voltage for the processor in standby modes	Ref
VCCST_PWRGD	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 must 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 the supplies come within specification. The signal must then transition monotonically to a high state.	Ref
VDDQ	Processor I/O supply voltage for DDR3.	Ref
VCOMP_OUT	Processor power reference for PEG/Display RCOMP.	Ref
VIDSOUT VIDSCLK VIDALERT#	VIDALERT#, VIDSCLK, and VIDSCLK comprise a three signal serial synchronous interface used to transfer power management information between the processor and the voltage regulator controllers.	Input GTL/ Output Open Drain Output Open Drain Input CMOS

6.12 Sense Signals


Table 6-14. Sense Signals

Signal Name	Description	Direction / Buffer Type
VCC_SENSE VSS_SENSE	VCC_SENSE and VSS_SENSE provide an isolated, low-impedance connection to the processor input V_{CC} voltage and ground. They can be used to sense or measure voltage near the silicon.	O A

6.13 Ground and Non-Critical to Function (NCTF) Signals

Table 6-15. Ground and Non-Critical to Function (NCTF) Signals

Signal Name	Description	Direction / Buffer Type
VSS	Processor ground node	GND
VSS_NCTF	Non-Critical to Function: These pins are for package mechanical reliability.	—



6.14 Processor Internal Pull-Up / Pull-Down Terminations

Table 6-16. Processor Internal Pull-Up / Pull-Down Terminations

Signal Name	Pull-Up / Pull-Down	Rail	Value
BPM[7:0]	Pull Up	VCCIO_TERM	40–60 Ω
PREQ#	Pull Up	VCCIO_TERM	40–60 Ω
TDI	Pull Up	VCCIO_TERM	30–70 Ω
TMS	Pull Up	VCCIO_TERM	30–70 Ω
CFG[17:0]	Pull Up	VCCIO_OUT	5–8 k Ω
CATERR#	Pull Up	VCCIO_TERM	30–70 Ω

§





7 Electrical Specifications

This chapter provides the processor electrical specifications including integrated voltage regulator (VR), V_{CC} Voltage Identification (VID), reserved and unused signals, signal groups, Test Access Points (TAP), DC specifications.

7.1 Integrated Voltage Regulator

A new feature to the processor is the integration of platform voltage regulators into the processor. Due to this integration, the processor has one main voltage rail (V_{CC}) and a voltage rail for the memory interface (V_{DDQ}) compared to six voltage rails on previous processors. The V_{CC} voltage rail will supply the integrated voltage regulators which in turn will regulate to the appropriate voltages for the cores, cache, system agent, and graphics. This integration allows the processor to better control on-die voltages to optimize between performance and power savings. The processor V_{CC} rail will remain a VID-based voltage with a loadline similar to the core voltage rail (also called V_{CC}) in previous processors.

7.2 Power and Ground Lands

The processor has VCC, VDDQ, and VSS (ground) lands for on-chip power distribution. All power lands must be connected to their respective processor power planes, while all VSS lands must be connected to the system ground plane. Use of multiple power and ground planes is recommended to reduce I^2R drop. The VCC lands must be supplied with the voltage determined by the processor **S**erial **V**oltage **I**dentification (SVID) interface. [Table 7-1, "VR 12.5 Voltage Identification"](#) specifies the voltage level for the various VIDs.

7.3 V_{CC} Voltage Identification (VID)

The processor uses three signals for the serial voltage identification interface to support automatic selection of voltages. The following table specifies the voltage level corresponding to the 8-bit VID value transmitted over serial VID. A '1' in this table refers to a high voltage level and a '0' refers to a low voltage level. If the voltage regulation circuit cannot supply the voltage that is requested, the voltage regulator must disable itself. VID signals are CMOS push/pull drivers. See [Table 7-7](#) for the DC specifications for these signals. The VID codes will change due to temperature and/or current load changes in order to minimize the power of the part. A voltage range is provided in [Section 7.8, "Voltage and Current Specifications"](#). The specifications are set so that one voltage regulator can operate with all supported frequencies.

Individual processor VID values may be set during manufacturing so that two devices at the same core frequency may have different default VID settings. This is shown in the VID range values in [Section 7.8, "Voltage and Current Specifications"](#). The processor provides the ability to operate while transitioning to an adjacent VID and its associated voltage. This will represent a DC shift in the loadline.



Table 7-1. VR 12.5 Voltage Identification (Sheet 1 of 7)

Bit7	Bit6	Bit5	Bit4	Bit3	Bit2	Bit1	Bit0	Hex	V _{CC}
0	0	0	0	0	0	0	0	00h	0.0000
0	0	0	0	0	0	0	1	01h	0.5000
0	0	0	0	0	0	1	0	02h	0.5100
0	0	0	0	0	0	1	1	03h	0.5200
0	0	0	0	0	1	0	0	04h	0.5300
0	0	0	0	0	1	0	1	05h	0.5400
0	0	0	0	0	1	1	0	06h	0.5500
0	0	0	0	0	1	1	1	07h	0.5600
0	0	0	0	1	0	0	0	08h	0.5700
0	0	0	0	1	0	0	1	09h	0.5800
0	0	0	0	1	0	1	0	0Ah	0.5900
0	0	0	0	1	0	1	1	0Bh	0.6000
0	0	0	0	1	1	0	0	0Ch	0.6100
0	0	0	0	1	1	0	1	0Dh	0.6200
0	0	0	0	1	1	1	0	0Eh	0.6300
0	0	0	0	1	1	1	1	0Fh	0.6400
0	0	0	1	0	0	0	0	10h	0.6500
0	0	0	1	0	0	0	1	11h	0.6600
0	0	0	1	0	0	1	0	12h	0.6700
0	0	0	1	0	0	1	1	13h	0.6800
0	0	0	1	0	1	0	0	14h	0.6900
0	0	0	1	0	1	0	1	15h	0.7000
0	0	0	1	0	1	1	0	16h	0.7100
0	0	0	1	0	1	1	1	17h	0.7200
0	0	0	1	1	0	0	0	18h	0.7300
0	0	0	1	1	0	0	1	19h	0.7400
0	0	0	1	1	0	1	0	1Ah	0.7500
0	0	0	1	1	0	1	1	1Bh	0.7600
0	0	0	1	1	1	0	0	1Ch	0.7700
0	0	0	1	1	1	0	1	1Dh	0.7800
0	0	0	1	1	1	1	0	1Eh	0.7900
0	0	0	1	1	1	1	1	1Fh	0.8000
0	0	1	0	0	0	0	0	20h	0.8100
0	0	1	0	0	0	0	1	21h	0.8200
0	0	1	0	0	0	1	0	22h	0.8300
0	0	1	0	0	0	1	1	23h	0.8400
0	0	1	0	0	1	0	0	24h	0.8500
0	0	1	0	0	1	0	1	25h	0.8600
0	0	1	0	0	1	1	0	26h	0.8700
0	0	1	0	0	1	1	1	27h	0.8800



Table 7-1. VR 12.5 Voltage Identification (Sheet 2 of 7)

Bit7	Bit6	Bit5	Bit4	Bit3	Bit2	Bit1	Bit0	Hex	V _{CC}
0	0	1	0	1	0	0	0	28h	0.8900
0	0	1	0	1	0	0	1	29h	0.9000
0	0	1	0	1	0	1	0	2Ah	0.9100
0	0	1	0	1	0	1	1	2Bh	0.9200
0	0	1	0	1	1	0	0	2Ch	0.9300
0	0	1	0	1	1	0	1	2Dh	0.9400
0	0	1	0	1	1	1	0	2Eh	0.9500
0	0	1	0	1	1	1	1	2Fh	0.9600
0	0	1	1	0	0	0	0	30h	0.9700
0	0	1	1	0	0	0	1	31h	0.9800
0	0	1	1	0	0	1	0	32h	0.9900
0	0	1	1	0	0	1	1	33h	1.0000
0	0	1	1	0	1	0	0	34h	1.0100
0	0	1	1	0	1	0	1	35h	1.0200
0	0	1	1	0	1	1	0	36h	1.0300
0	0	1	1	0	1	1	1	37h	1.0400
0	0	1	1	1	0	0	0	38h	1.0500
0	0	1	1	1	0	0	1	39h	1.0600
0	0	1	1	1	0	1	0	3Ah	1.0700
0	0	1	1	1	0	1	1	3Bh	1.0800
0	0	1	1	1	1	0	0	3Ch	1.0900
0	0	1	1	1	1	0	1	3Dh	1.1000
0	0	1	1	1	1	1	0	3Eh	1.1100
0	0	1	1	1	1	1	1	3Fh	1.1200
0	1	0	0	0	0	0	0	40h	1.1300
0	1	0	0	0	0	0	1	41h	1.1400
0	1	0	0	0	0	1	0	42h	1.1500
0	1	0	0	0	0	1	1	43h	1.1600
0	1	0	0	0	1	0	0	44h	1.1700
0	1	0	0	0	1	0	1	45h	1.1800
0	1	0	0	0	1	1	0	46h	1.1900
0	1	0	0	0	1	1	1	47h	1.2000
0	1	0	0	1	0	0	0	48h	1.2100
0	1	0	0	1	0	0	1	49h	1.2200
0	1	0	0	1	0	1	0	4Ah	1.2300
0	1	0	0	1	0	1	1	4Bh	1.2400
0	1	0	0	1	1	0	0	4Ch	1.2500
0	1	0	0	1	1	0	1	4Dh	1.2600
0	1	0	0	1	1	1	0	4Eh	1.2700
0	1	0	0	1	1	1	1	4Fh	1.2800



Table 7-1. VR 12.5 Voltage Identification (Sheet 3 of 7)

Bit7	Bit6	Bit5	Bit4	Bit3	Bit2	Bit1	Bit0	Hex	V _{CC}
0	1	0	1	0	0	0	0	50h	1.2900
0	1	0	1	0	0	0	1	51h	1.3000
0	1	0	1	0	0	1	0	52h	1.3100
0	1	0	1	0	0	1	1	53h	1.3200
0	1	0	1	0	1	0	0	54h	1.3300
0	1	0	1	0	1	0	1	55h	1.3400
0	1	0	1	0	1	1	0	56h	1.3500
0	1	0	1	0	1	1	1	57h	1.3600
0	1	0	1	1	0	0	0	58h	1.3700
0	1	0	1	1	0	0	1	59h	1.3800
0	1	0	1	1	0	1	0	5Ah	1.3900
0	1	0	1	1	0	1	1	5Bh	1.4000
0	1	0	1	1	1	0	0	5Ch	1.4100
0	1	0	1	1	1	0	1	5Dh	1.4200
0	1	0	1	1	1	1	0	5Eh	1.4300
0	1	0	1	1	1	1	1	5Fh	1.4400
0	1	1	0	0	0	0	0	60h	1.4500
0	1	1	0	0	0	0	1	61h	1.4600
0	1	1	0	0	0	1	0	62h	1.4700
0	1	1	0	0	0	1	1	63h	1.4800
0	1	1	0	0	1	0	0	64h	1.4900
0	1	1	0	0	1	0	1	65h	1.5000
0	1	1	0	0	1	1	0	66h	1.5100
0	1	1	0	0	1	1	1	67h	1.5200
0	1	1	0	1	0	0	0	68h	1.5300
0	1	1	0	1	0	0	1	69h	1.5400
0	1	1	0	1	0	1	0	6Ah	1.5500
0	1	1	0	1	0	1	1	6Bh	1.5600
0	1	1	0	1	1	0	0	6Ch	1.5700
0	1	1	0	1	1	0	1	6Dh	1.5800
0	1	1	0	1	1	1	0	6Eh	1.5900
0	1	1	0	1	1	1	1	6Fh	1.6000
0	1	1	1	0	0	0	0	70h	1.6100
0	1	1	1	0	0	0	1	71h	1.6200
0	1	1	1	0	0	1	0	72h	1.6300
0	1	1	1	0	0	1	1	73h	1.6400
0	1	1	1	0	1	0	0	74h	1.6500
0	1	1	1	0	1	0	1	75h	1.6600
0	1	1	1	0	1	1	0	76h	1.6700
0	1	1	1	0	1	1	1	77h	1.6800



Table 7-1. VR 12.5 Voltage Identification (Sheet 4 of 7)

Bit7	Bit6	Bit5	Bit4	Bit3	Bit2	Bit1	Bit0	Hex	V _{CC}
0	1	1	1	1	0	0	0	78h	1.6900
0	1	1	1	1	0	0	1	79h	1.7000
0	1	1	1	1	0	1	0	7Ah	1.7100
0	1	1	1	1	0	1	1	7Bh	1.7200
0	1	1	1	1	1	0	0	7Ch	1.7300
0	1	1	1	1	1	0	1	7Dh	1.7400
0	1	1	1	1	1	1	0	7Eh	1.7500
0	1	1	1	1	1	1	1	7Fh	1.7600
1	0	0	0	0	0	0	0	80h	1.7700
1	0	0	0	0	0	0	1	81h	1.7800
1	0	0	0	0	0	1	0	82h	1.7900
1	0	0	0	0	0	1	1	83h	1.8000
1	0	0	0	0	1	0	0	84h	1.8100
1	0	0	0	0	1	0	1	85h	1.8200
1	0	0	0	0	1	1	0	86h	1.8300
1	0	0	0	0	1	1	1	87h	1.8400
1	0	0	0	1	0	0	0	88h	1.8500
1	0	0	0	1	0	0	1	89h	1.8600
1	0	0	0	1	0	1	0	8Ah	1.8700
1	0	0	0	1	0	1	1	8Bh	1.8800
1	0	0	0	1	1	0	0	8Ch	1.8900
1	0	0	0	1	1	0	1	8Dh	1.9000
1	0	0	0	1	1	1	0	8Eh	1.9100
1	0	0	0	1	1	1	1	8Fh	1.9200
1	0	0	1	0	0	0	0	90h	1.9300
1	0	0	1	0	0	0	1	91h	1.9400
1	0	0	1	0	0	1	0	92h	1.9500
1	0	0	1	0	0	1	1	93h	1.9600
1	0	0	1	0	1	0	0	94h	1.9700
1	0	0	1	0	1	0	1	95h	1.9800
1	0	0	1	0	1	1	0	96h	1.9900
1	0	0	1	0	1	1	1	97h	2.0000
1	0	0	1	1	0	0	0	98h	2.0100
1	0	0	1	1	0	0	1	99h	2.0200
1	0	0	1	1	0	1	0	9Ah	2.0300
1	0	0	1	1	0	1	1	9Bh	2.0400
1	0	0	1	1	1	0	0	9Ch	2.0500
1	0	0	1	1	1	0	1	9Dh	2.0600
1	0	0	1	1	1	1	0	9Eh	2.0700
1	0	0	1	1	1	1	1	9Fh	2.0800



Table 7-1. VR 12.5 Voltage Identification (Sheet 5 of 7)

Bit7	Bit6	Bit5	Bit4	Bit3	Bit2	Bit1	Bit0	Hex	V _{CC}
1	0	1	0	0	0	0	0	A0h	2.0900
1	0	1	0	0	0	0	1	A1h	2.1000
1	0	1	0	0	0	1	0	A2h	2.1100
1	0	1	0	0	0	1	1	A3h	2.1200
1	0	1	0	0	1	0	0	A4h	2.1300
1	0	1	0	0	1	0	1	A5h	2.1400
1	0	1	0	0	1	1	0	A6h	2.1500
1	0	1	0	0	1	1	1	A7h	2.1600
1	0	1	0	1	0	0	0	A8h	2.1700
1	0	1	0	1	0	0	1	A9h	2.1800
1	0	1	0	1	0	1	0	AAh	2.1900
1	0	1	0	1	0	1	1	ABh	2.2000
1	0	1	0	1	1	0	0	ACH	2.2100
1	0	1	0	1	1	0	1	ADh	2.2200
1	0	1	0	1	1	1	0	A Eh	2.2300
1	0	1	0	1	1	1	1	AFh	2.2400
1	0	1	1	0	0	0	0	B0h	2.2500
1	0	1	1	0	0	0	1	B1h	2.2600
1	0	1	1	0	0	1	0	B2h	2.2700
1	0	1	1	0	0	1	1	B3h	2.2800
1	0	1	1	0	1	0	0	B4h	2.2900
1	0	1	1	0	1	0	1	B5h	2.3000
1	0	1	1	0	1	1	0	B6h	2.3100
1	0	1	1	0	1	1	1	B7h	2.3200
1	0	1	1	1	0	0	0	B8h	2.3300
1	0	1	1	1	0	0	1	B9h	2.3400
1	0	1	1	1	0	1	0	BAh	2.3500
1	0	1	1	1	0	1	1	BBh	2.3600
1	0	1	1	1	1	0	0	BCh	2.3700
1	0	1	1	1	1	0	1	BDh	2.3800
1	0	1	1	1	1	1	0	BEh	2.3900
1	0	1	1	1	1	1	1	BFh	2.4000
1	1	0	0	0	0	0	0	C0h	2.4100
1	1	0	0	0	0	0	1	C1h	2.4200
1	1	0	0	0	0	1	0	C2h	2.4300
1	1	0	0	0	0	1	1	C3h	2.4400
1	1	0	0	0	1	0	0	C4h	2.4500
1	1	0	0	0	1	0	1	C5h	2.4600
1	1	0	0	0	1	1	0	C6h	2.4700
1	1	0	0	0	1	1	1	C7h	2.4800



Table 7-1. VR 12.5 Voltage Identification (Sheet 6 of 7)

Bit7	Bit6	Bit5	Bit4	Bit3	Bit2	Bit1	Bit0	Hex	V _{CC}
1	1	0	0	1	0	0	0	C8h	2.4900
1	1	0	0	1	0	0	1	C9h	2.5000
1	1	0	0	1	0	1	0	CAh	2.5100
1	1	0	0	1	0	1	1	CBh	2.5200
1	1	0	0	1	1	0	0	CCh	2.5300
1	1	0	0	1	1	0	1	CDh	2.5400
1	1	0	0	1	1	1	0	CEh	2.5500
1	1	0	0	1	1	1	1	CFh	2.5600
1	1	0	1	0	0	0	0	D0h	2.5700
1	1	0	1	0	0	0	1	D1h	2.5800
1	1	0	1	0	0	1	0	D2h	2.5900
1	1	0	1	0	0	1	1	D3h	2.6000
1	1	0	1	0	1	0	0	D4h	2.6100
1	1	0	1	0	1	0	1	D5h	2.6200
1	1	0	1	0	1	1	0	D6h	2.6300
1	1	0	1	0	1	1	1	D7h	2.6400
1	1	0	1	1	0	0	0	D8h	2.6500
1	1	0	1	1	0	0	1	D9h	2.6600
1	1	0	1	1	0	1	0	DAh	2.6700
1	1	0	1	1	0	1	1	DBh	2.6800
1	1	0	1	1	1	0	0	DCh	2.6900
1	1	0	1	1	1	0	1	DDh	2.7000
1	1	0	1	1	1	1	0	DEh	2.7100
1	1	0	1	1	1	1	1	DFh	2.7200
1	1	1	0	0	0	0	0	E0h	2.7300
1	1	1	0	0	0	0	1	E1h	2.7400
1	1	1	0	0	0	1	0	E2h	2.7500
1	1	1	0	0	0	1	1	E3h	2.7600
1	1	1	0	0	1	0	0	E4h	2.7700
1	1	1	0	0	1	0	1	E5h	2.7800
1	1	1	0	0	1	1	0	E6h	2.7900
1	1	1	0	0	1	1	1	E7h	2.8000
1	1	1	0	1	0	0	0	E8h	2.8100
1	1	1	0	1	0	0	1	E9h	2.8200
1	1	1	0	1	0	1	0	EAh	2.8300
1	1	1	0	1	0	1	1	EBh	2.8400
1	1	1	0	1	1	0	0	ECh	2.8500
1	1	1	0	1	1	0	1	EDh	2.8600
1	1	1	0	1	1	1	0	EEh	2.8700
1	1	1	0	1	1	1	1	EFh	2.8800

Table 7-1. VR 12.5 Voltage Identification (Sheet 7 of 7)

Bit7	Bit6	Bit5	Bit4	Bit3	Bit2	Bit1	Bit0	Hex	V _{CC}
1	1	1	1	0	0	0	0	F0h	2.8900
1	1	1	1	0	0	0	1	F1h	2.9000
1	1	1	1	0	0	1	0	F2h	2.9100
1	1	1	1	0	0	1	1	F3h	2.9200
1	1	1	1	0	1	0	0	F4h	2.9300
1	1	1	1	0	1	0	1	F5h	2.9400
1	1	1	1	0	1	1	0	F6h	2.9500
1	1	1	1	0	1	1	1	F7h	2.9600
1	1	1	1	1	0	0	0	F8h	2.9700
1	1	1	1	1	0	0	1	F9h	2.9800
1	1	1	1	1	0	1	0	FAh	2.9900
1	1	1	1	1	0	1	1	FBh	3.0000
1	1	1	1	1	1	0	0	FCh	3.0100
1	1	1	1	1	1	0	1	FDh	3.0200
1	1	1	1	1	1	1	0	FEh	3.0300
1	1	1	1	1	1	1	1	FFh	3.0400

7.4 Reserved or Unused 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
- RSVD_NCTF – these signals are non-critical to function and may be left unconnected

Arbitrary connection of these signals to V_{CC}, V_{DDQ}, V_{SS}, or to any other signal (including each other) may result in component malfunction or incompatibility with future processors. See [Chapter 6, “Signal Description”](#) for a pin listing of the processor and the location of all reserved 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 must be used when tying bi-directional signals to power or ground. When tying any signal to power or ground, a resistor will also allow for system testability.



7.5 Signal Groups

Signals are grouped by buffer type and similar characteristics as listed in the following table. The buffer type indicates which signaling technology and specifications apply to the signals. All the differential signals and selected DDR3L and Control Sideband signals have On-Die Termination (ODT) resistors. Some signals do not have ODT and need to be terminated on the board.

Note: All Control Sideband Asynchronous signals are required to be asserted/de-asserted for at least 10 BCLKs with maximum Trise/Tfall of 6 ns in order for the processor to recognize the proper signal state. See [Section 7.7, "DC Specifications"](#).

Table 7-2. Signal Groups (Sheet 1 of 2)

Signal Group	Type	Signals
System Reference Clock		
Differential	CMOS Input	BCLKP, BCLKN, DPLL_REF_CLKP, DPLL_REF_CLKN, SSC_DPLL_REF_CLKP, SSC_DPLL_REF_CLKN
DDR3L Reference Clocks ²		
Differential	DDR3L Output	SA_CKP[3:0], SA_CKN[3:0], SB_CKP[3:0], SB_CKN[3:0]
DDR3L Command Signals ²		
Single ended	DDR3L Output	SA_BS[2:0], SB_BS[2:0], SA_WE#, SB_WE#, SA_RAS#, SB_RAS#, SA_CAS#, SB_CAS#, SA_MA[15:0], SB_MA[15:0]
DDR3L Control Signals ²		
Single ended	DDR3L Output	SA_CKE[3:0], SB_CKE[3:0], SA_CS#[3:0], SB_CS#[3:0], SA_ODT[3:0], SB_ODT[3:0]
Single ended	CMOS Output	SM_DRAMRST#
DDR3L Data Signals ²		
Single ended	DDR3L Bi-directional	SA_DQ[63:0], SB_DQ[63:0]
Differential	DDR3L Bi-directional	SA_DQSP[7:0], SA_DQSN[7:0], SB_DQSP[7:0], SB_DQSN[7:0]
DDR3L Compensation		
	Analog Input	SM_RCOMP[2:0]
DDR3L Reference Voltage Signals		
	DDR3L Output	SM_VREF, SA_DIMM_VREFDQ, SB_DIMM_VREFDQ
Testability (ITP/XDP)		
Single ended	CMOS Input	TCK, TDI, TMS, TRST#
Single ended	GTL	TDO
Single ended	Output	DBR#
Single ended	GTL	BPM#[7:0]
Single ended	GTL	PREQ#
Single ended	GTL	PRDY#
Control Sideband		
Single ended	GTL Input/Open Drain Output	PROCHOT#



Table 7-2. Signal Groups (Sheet 2 of 2)

Signal Group	Type	Signals
Single ended	Asynchronous CMOS Output	THERMTRIP#, IVR_ERROR
Single ended	GTL	CATERR#
Single ended	Asynchronous CMOS Input	PM_SYNC, RESET#, PWRGOOD, PWR_DEBUG#
Single ended	Asynchronous Bi-directional	PECI
Single ended	GTL Bi-directional	CFG[19:0]
Single ended	Analog Input	SM_RCOMP[2:0]
Voltage Regulator		
Single ended	CMOS Input	VR_READY
Single ended	CMOS Input	VIDALERT#
Single ended	Open Drain Output	VIDSCLK
Single ended	GTL Input/Open Drain Output	VIDSOUT
Differential	Analog Output	VCC_SENSE, VSS_SENSE
Power / Ground / Other		
Single ended	Power	VCC, VDDQ, VCCST
	Ground	VSS, VSS_NCTF ³
	No Connect	RSVD, RSVD_NCTF
	Test Point	RSVD_TP
	Other	SKTOCC#,
PCI Express* Graphics		
Differential	PCI Express Input	PEG_RXP[15:0], PEG_RXN[15:0]
Differential	PCI Express Output	PEG_TXP[15:0], PEG_TXN[15:0]
Single ended	Analog Input	PEG_RCOMP
Digital Media Interface (DMI)		
Differential	DMI Input	DMI_RXP[3:0], DMI_RXN[3:0]
Differential	DMI Output	DMI_TXP[3:0], DMI_TXN[3:0]
Digital Display Interface		
Differential	DDI Output	DDIB_TXP[3:0], DDIB_TXN[3:0], DDIC_TXP[3:0], DDIC_TXN[3:0], DDID_TXP[3:0], DDID_TXN[3:0]
Intel® FDI		
Single ended	CMOS Input	FDI_CS SYNC
Single ended	Asynchronous CMOS Input	DISP_INT
Differential	FDI Output	FDI_TXP[1:0], FDI_TXN[1:0]
Notes:		
1. See Chapter 6, "Signal Description" for signal description details.		
2. SA and SB refer to DDR3L Channel A and DDR3L Channel B.		



7.6 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. A few of the I/O pins may support only one of those standards.

7.7 DC Specifications

The processor DC specifications in this section are defined at the processor pins, unless noted otherwise. See [Chapter 6, "Signal Description"](#) for the processor pin listings and signal definitions.

- The DC specifications for the DDR3L signals are listed in the [Section 7.8, "Voltage and Current Specifications"](#).
- [Section 7.8, "Voltage and Current Specifications"](#) 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 DC rails include dynamic load currents at switching frequencies up to 1 MHz.

7.8 Voltage and Current Specifications

Table 7-3. Processor Core Active and Idle Mode DC Voltage and Current Specifications (Sheet 1 of 2)

Symbol	Parameter	Min	Typ	Max	Unit	Note ¹
Operational VID	VID Range	1.65	2013D: 1.75 2013C: 1.75 2013B: 1.75 2013A: 1.75	1.86	V	2
Idle VID (package C6)	VID Range	1.5	1.6	1.65	V	2
R_DC_LL	Loadline slope within the VR regulation loop capability		2013D PCG: -1.5 2013C PCG: -1.5 2013B PCG: -1.5 2013A PCG: -1.5		mΩ	3, 5, 6, 8
R_AC_LL	Loadline slope in response to dynamic load increase events		2013D PCG: -2.4 2013C PCG: -2.4 2013B PCG: -2.4 2013A PCG: -2.4		mΩ	—
R_AC_LL_OS	Loadline slope in response to dynamic load release events		2013D PCG: -3.0 2013C PCG: -3.0 2013B PCG: -3.0 2013A PCG: -3.0		mΩ	—
T_OVS	Overshoot time			500	μs	
V_OVS	Overshoot			50	mV	
V _{CC} TOB	V _{CC} Tolerance Band	± 20 (PS0, PS1, PS2, PS3)			mV	3, 5, 6, 7, 8



Table 7-3. Processor Core Active and Idle Mode DC Voltage and Current Specifications (Sheet 2 of 2)

Symbol	Parameter	Min	Typ	Max	Unit	Note ¹
V _{CC} Ripple	Ripple		± 10 (PS0) ± 15 (PS1) +50/-15 (PS2) +60/-15 (PS3)		mV	3, 5, 6, 7, 8
V _{CC,BOOT}	Default V _{CC} voltage for initial power up	—	1.70	—	V	—
I _{CC}	2013D PCG I _{CC}	—	—	95	A	4, 8
I _{CC}	2013C PCG I _{CC}	—	—	75	A	4, 8
I _{CC}	2013B PCG I _{CC}	—	—	58	A	4, 8
I _{CC}	2013A PCG I _{CC}	—	—	48	A	4, 8
P _{MAX}	2013D PCG P _{MAX}	—	—	153	W	9
P _{MAX}	2013C PCG P _{MAX}	—	—	121	W	9
P _{MAX}	2013B PCG P _{MAX}	—	—	99	W	9
P _{MAX}	2013A PCG P _{MAX}	—	—	83	W	9

Notes:

- 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.
- 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. 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).
- The voltage specification requirements are measured across V_{CC_SENSE} and V_{VSS_SENSE} lands at the socket with a 20-MHz bandwidth oscilloscope, 1.5 pF maximum probe capacitance, and 1-MΩ 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.
- I_{CC_MAX} specification is based on the V_{CC} loadline at worst case (highest) tolerance and ripple.
- The V_{CC} specifications represent static and transient limits.
- The loadlines specify voltage limits at the die measured at the V_{CC_SENSE} and V_{VSS_SENSE} lands. Voltage regulation feedback for voltage regulator circuits must also be taken from processor V_{CC_SENSE} and V_{VSS_SENSE} lands.
- PS_x refers to the voltage regulator power state as set by the SVID protocol.
- PCG is Platform Compatibility Guide (previously known as FMB). These guidelines are for estimation purposes only.
- P_{MAX} is the maximum power the processor will dissipate as measured at V_{CC_SENSE} and V_{VSS_SENSE} lands. The processor may draw this power for up to 10 ms before it regulates to PL2.

Table 7-4. Memory Controller (V_{DDQ}) Supply DC Voltage and Current Specifications (Sheet 1 of 2)

Symbol	Parameter	Min	Typ	Max	Unit	Note	
V _{DDQ} (DC+AC) DDR3L	Processor I/O supply voltage for DDR3L (DC + AC specification)		Typ-5%	1.5	Typ+5%	V	2, 3
V _{DDQ} (DC+AC) DDR3L	Processor I/O supply voltage for DDR3L (DC + AC specification)		Typ-5%	1.35	Typ+5%	V	2, 3, 5


Table 7-4. Memory Controller (V_{DDQ}) Supply DC Voltage and Current Specifications (Sheet 2 of 2)

Symbol	Parameter	Min	Typ	Max	Unit	Note
I_{CCMAX_VDDQ} (DDR3L)	Max Current for V_{DDQ} Rail	—	—	2.5	A	1
I_{CCAVG_VDDQ} (Standby)	Average Current for V_{DDQ} Rail during Standby	—	12	20	mA	4
Notes: <ol style="list-style-type: none"> The current supplied to the SO-DIMM modules is not included in this specification. Includes AC and DC error, where the AC noise is bandwidth limited to under 20 MHz. No requirement on the breakdown of AC versus DC noise. Measured at 50 °C This specification applies to UP Server/workstation processors paired with a PCH configured with Intel Server Platform Services FW. 						

Table 7-5. V_{CCIO_OUT} , V_{COMP_OUT} , and V_{CCIO_TERM}

Symbol	Parameter	Typ	Max	Units	Notes
V_{CCIO_OUT}	Termination Voltage	1.0	—	V	3, 4
I_{CCIO_OUT}	Maximum External Load	—	300	mA	3, 4
V_{COMP_OUT}	Termination Voltage	1.0	—	V	1, 3, 4
V_{CCIO_TERM}	Termination Voltage	1.0	—	V	2
Notes: <ol style="list-style-type: none"> V_{COMP_OUT} may only be used to connect to PEG_RCOMP and DP_RCOMP. Internal processor power for signal termination. N/A N/A 					

Table 7-6. DDR3L Signal Group DC Specifications (Sheet 1 of 2)

Symbol	Parameter	Min	Typ	Max	Units	Notes ¹
V_{IL}	Input Low Voltage	—	$V_{DDQ}/2$	$0.43 \cdot V_{DDQ}$	V	2, 4, 11
V_{IH}	Input High Voltage	$0.57 \cdot V_{DDQ}$	$V_{DDQ}/2$	—	V	3, 11
V_{IL}	Input Low Voltage (SM_DRAMPWROK)	—	—	$0.15 \cdot V_{DDQ}$	V	—
V_{IH}	Input High Voltage (SM_DRAMPWROK)	$0.45 \cdot V_{DDQ}$	—	1.0	V	10, 12
$R_{ON_UP}(DQ)$	DDR3L Data Buffer pull-up Resistance	20	26	32	Ω	5, 11
$R_{ON_DN}(DQ)$	DDR3L Data Buffer pull-down Resistance	20	26	32	Ω	5, 11
$R_{ODT}(DQ)$	DDR3L On-die termination equivalent resistance for data signals	38	50	62	Ω	11
$V_{ODT}(DC)$	DDR3L On-die termination DC working point (driver set to receive mode)	$0.45 \cdot V_{DDQ}$	$0.5 \cdot V_{DDQ}$	$0.55 \cdot V_{DDQ}$	V	11
$R_{ON_UP}(CK)$	DDR3L Clock Buffer pull-up Resistance	20	26	32	Ω	5, 11, 13
$R_{ON_DN}(CK)$	DDR3L Clock Buffer pull-down Resistance	20	26	32	Ω	5, 11, 13
$R_{ON_UP}(CMD)$	DDR3L Command Buffer pull-up Resistance	15	20	25	Ω	5, 11, 13



Table 7-6. DDR3L Signal Group DC Specifications (Sheet 2 of 2)

Symbol	Parameter	Min	Typ	Max	Units	Notes ¹
R _{ON_DN} (CMD)	DDR3L Command Buffer pull-down Resistance	15	20	25	Ω	5, 11, 13
R _{ON_UP} (CTL)	DDR3L Control Buffer pull-up Resistance	19	25	31	Ω	5, 11, 13
R _{ON_DN} (CTL)	DDR3L Control Buffer pull-down Resistance	19	25	31	Ω	5, 11, 13
R _{ON_UP} (RST)	DDR3L Reset Buffer pull-up Resistance	40	80	130	Ω	—
R _{ON_DN} (RST)	DDR3L Reset Buffer pull-up Resistance	40	80	130	Ω	—
I _{LI}	Input Leakage Current (DQ, CK) 0V 0.2*V _{DDQ} 0.8*V _{DDQ}	—	—	0.7	mA	—
I _{LI}	Input Leakage Current (CMD, CTL) 0V 0.2*V _{DDQ} 0.8*V _{DDQ}	—	—	1.0	mA	—
SM_RCOMP0	Command COMP Resistance	99	100	101	Ω	8
SM_RCOMP1	Data COMP Resistance	74.25	75	75.75	Ω	8
SM_RCOMP2	ODT COMP Resistance	99	100	101	Ω	8
<p>Notes:</p> <ol style="list-style-type: none"> Unless otherwise noted, all specifications in this table apply to all processor frequencies. V_{IL} is defined as the maximum voltage level at a receiving agent that will be interpreted as a logical low value. V_{IH} is defined as the minimum voltage level at a receiving agent that will be interpreted as a logical high value. V_{IH} and V_{OH} may experience excursions above V_{DDQ}. However, input signal drivers must comply with the signal quality specifications. This is the pull up/down driver resistance. R_{TERM} is the termination on the DIMM and is not controlled by the processor. The minimum and maximum values for these signals are programmable by BIOS to one of the two sets. SM_RCOMPx resistance must be provided on the system board with 1% resistors. SM_RCOMPx resistors are to V_{SS}. DDR3L values are pre-silicon estimations and are subject to change. SM_DRAMPWROK rise and fall time must be < 50 ns measured between V_{DDQ} *0.15 and V_{DDQ} *0.47. SM_VREF is defined as V_{DDQ}/2. Maximum-minimum range is correct but center point is subject to change during MRC boot training. Processor may be damaged if V_{IH} exceeds the maximum voltage for extended periods. The MRC during boot training might optimize R_{ON} outside the range specified. 						

Table 7-7. Digital Display Interface Group DC Specifications

Symbol	Parameter	Min	Typ	Max	Units
V _{IL}	HPD Input Low Voltage	—	—	0.8	V
V _{IH}	HPD Input High Voltage	2.25	—	3.6	V
Vaux(Tx)	Aux peak-to-peak voltage at transmitting device	0.39	—	1.38	V
Vaux(Rx)	Aux peak-to-peak voltage at receiving device	0.32	—	1.36	V


Table 7-8. Embedded DisplayPort* (eDP) Group DC Specifications

Symbol	Parameter	Min	Typ	Max	Units
V _{IL}	HPD Input Low Voltage	0.02	—	0.21	V
V _{IH}	HPD Input High Voltage	0.84	—	1.05	V
V _{OL}	eDP_DISP_UTIL Output Low Voltage	0.1*V _{CC}	—	—	V
V _{OH}	eDP_DISP_UTIL Output High Voltage	0.9*V _{CC}	—	—	V
R _{UP}	eDP_DISP_UTIL Internal pull-up	100	—	—	Ω
R _{DOWN}	eDP_DISP_UTIL Internal pull-down	100	—	—	Ω
V _{aux} (Tx)	Aux peak-to-peak voltage at transmitting device	0.39	—	1.38	V
V _{aux} (Rx)	Aux peak-to-peak voltage at receiving device	0.32	—	1.36	V
eDP_RCOMP DP_RCOMP	COMP Resistance	24.75	25	25.25	Ω

Note:
1. COMP resistance is to VCOMP_OUT.

Table 7-9. CMOS Signal Group DC Specifications

Symbol	Parameter	Min	Max	Units	Notes ¹
V _{IL}	Input Low Voltage	—	V _{CCIO_OUT} * 0.3	V	2, 5
V _{IH}	Input High Voltage	V _{CCIO_OUT} * 0.7	—	V	2, 4, 5
V _{OL}	Output Low Voltage	—	V _{CCIO_OUT} * 0.1	V	2
V _{OH}	Output High Voltage	V _{CCIO_OUT} * 0.9	—	V	2, 4
R _{ON}	Buffer on Resistance	23	73	Ω	-
I _{LI}	Input Leakage Current	—	±150	μA	3

Notes:
1. Unless otherwise noted, all specifications in this table apply to all processor frequencies.
2. The V_{CCIO_OUT} referred to in these specifications refers to instantaneous V_{CCIO_OUT}.
3. For VIN between "0" V and V_{CCIO_OUT}. Measured when the driver is tri-stated.
4. V_{IH} and V_{OH} may experience excursions above V_{CCIO_OUT}. However, input signal drivers must comply with the signal quality specifications.
5. N/A

Table 7-10. GTL Signal Group and Open Drain Signal Group DC Specifications (Sheet 1 of 2)

Symbol	Parameter	Min	Max	Units	Notes ¹
V _{IL}	Input Low Voltage (TAP, except TCK)	—	V _{CCIO_TERM} * 0.6	V	2, 5
V _{IH}	Input High Voltage (TAP, except TCK)	V _{CCIO_TERM} * 0.72	—	V	2, 4, 5
V _{IL}	Input Low Voltage (TCK)	—	V _{CCIO_TERM} * 0.4	V	2, 5
V _{IH}	Input High Voltage (TCK)	V _{CCIO_TERM} * 0.8	—	V	2, 4, 5
V _{HYSTERESIS}	Hysteresis Voltage	V _{CCIO_TERM} * 0.2	—	V	—
R _{ON}	Buffer on Resistance (TDO)	12	28	Ω	—
V _{IL}	Input Low Voltage (other GTL)	—	V _{CCIO_TERM} * 0.6	V	2, 5
V _{IH}	Input High Voltage (other GTL)	V _{CCIO_TERM} * 0.72	—	V	2, 4, 5
R _{ON}	Buffer on Resistance (CFG/BPM)	16	24	Ω	—



Table 7-10. GTL Signal Group and Open Drain Signal Group DC Specifications (Sheet 2 of 2)

Symbol	Parameter	Min	Max	Units	Notes ¹
R _{ON}	Buffer on Resistance (other GTL)	12	28	Ω	—
I _{LI}	Input Leakage Current	—	±150	μA	3
Notes: 1. Unless otherwise noted, all specifications in this table apply to all processor frequencies. 2. The V _{CCIO_OUT} referred to in these specifications refers to instantaneous V _{CCIO_OUT} . 3. For V _{IN} between 0 V and V _{CCIO_TERM} . Measured when the driver is tri-stated. 4. V _{IH} and V _{OH} may experience excursions above V _{CCIO_TERM} . However, input signal drivers must comply with the signal quality specifications. 5. N/A					

Table 7-11. PCI Express* DC Specifications

Symbol	Parameter	Min	Typ	Max	Units	Notes ¹
Z _{TX-DIFF-DC}	DC Differential Tx Impedance (Gen 1 Only)	80	—	120	Ω	1, 6
Z _{TX-DIFF-DC}	DC Differential Tx Impedance (Gen 2 and Gen 3)	—	—	120	Ω	1, 6
Z _{RX-DC}	DC Common Mode Rx Impedance	40	—	60	Ω	1, 4, 5
Z _{RX-DIFF-DC}	DC Differential Rx Impedance (Gen1 Only)	80	—	120	Ω	1
PEG_RCOMP	Comp Resistance	24.75	25	25.25	Ω	2, 3
Notes: 1. See the <i>PCI Express Base Specification</i> for more details. 2. PEG_RCOMP should be connected to V _{COMP_OUT} through a 25 Ω ±1% resistor. 3. Intel allows using 24.9 Ω ±1% resistors. 4. DC impedance limits are needed to ensure Receiver detect. 5. The Rx DC Common Mode Impedance must be present when the Receiver terminations are first enabled to ensure that the Receiver Detect occurs properly. Compensation of this impedance can start immediately and the 15 Rx Common Mode Impedance (constrained by RLRX-CM to 50 Ω ±20%) must be within the specified range by the time Detect is entered. 6. Low impedance defined during signaling. Parameter is captured for 5.0 GHz by RLTX-DIFF.						

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

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

Table 7-12. PECCI DC Electrical Limits (Sheet 1 of 2)

Symbol	Definition and Conditions	Min	Max	Units	Notes ¹
R _{up}	Internal pull up resistance	15	45	Ω	3
V _{in}	Input Voltage Range	-0.15	V _{CCIO_TERM} + 0.15	V	—
V _{hysteresis}	Hysteresis	0.1 * V _{CCIO_TERM}	N/A	V	—
V _n	Negative-Edge Threshold Voltage	0.275 * V _{CCIO_TERM}	0.500 * V _{CCIO_TERM}	V	—
V _p	Positive-Edge Threshold Voltage	0.550 * V _{CCIO_TERM}	0.725 * V _{CCIO_TERM}	V	—
C _{bus}	Bus Capacitance per Node	N/A	10	pF	—
C _{pad}	Pad Capacitance	0.7	1.8	pF	—



Table 7-12. PECE DC Electrical Limits (Sheet 2 of 2)

Symbol	Definition and Conditions	Min	Max	Units	Notes ¹
Ileak000	leakage current at 0 V	—	0.6	mA	—
Ileak025	leakage current at 0.25* V _{CCIO_TERM}	—	0.4	mA	—
Ileak050	leakage current at 0.50* V _{CCIO_TERM}	—	0.2	mA	—
Ileak075	leakage current at 0.75* V _{CCIO_TERM}	—	0.13	mA	—
Ileak100	leakage current at V _{CCIO_TERM}	—	0.10	mA	—

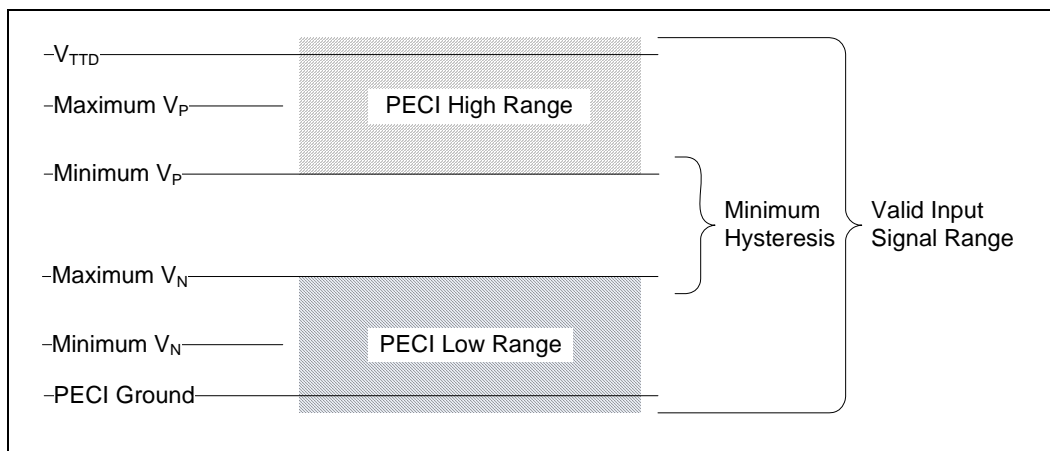
Notes:

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

7.8.2 Input Device Hysteresis

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

Figure 7-1. Input Device Hysteresis



§






8 Package Mechanical Specifications

The processor is packaged in a Flip-Chip Land Grid Array package that interfaces with the motherboard using the LGA1150 socket. The package consists of a processor mounted on a substrate land-carrier. An integrated heat spreader (IHS) is attached to the package substrate and core and serves as the mating surface for processor thermal solutions, such as a heatsink. The following figure shows a sketch of the processor package components and how they are assembled together.

The package components shown in the following figure include the following:

1. Integrated Heat Spreader (IHS)
2. Thermal Interface Material (TIM)
3. Processor core (die)
4. Package substrate
5. Capacitors

Figure 8-1. Processor Package Assembly Sketch



8.1 Processor Component Keep-Out Zone

The processor may contain components on the substrate that define component keep-out zone requirements. A thermal and mechanical solution design must not intrude into the required keep-out zones. Decoupling capacitors are typically mounted to the land-side of the package substrate. Refer to the *LGA1150 Socket Application Guide* for keep-out zones. The location and quantity of package capacitors may change due to manufacturing efficiencies but will remain within the component keep-in. This keep-in zone includes solder paste and is a post reflow maximum height for the components.

8.2 Package Loading Specifications

The following table provides dynamic and static load specifications for the processor package. These mechanical maximum load limits should not be exceeded during heatsink assembly, shipping conditions, or standard use condition. Also, any



mechanical system or component testing should not exceed the maximum limits. The processor package substrate should not be used as a mechanical reference or load-bearing surface for thermal and mechanical solution.

Figure 8-2. Processor Loading Specifications

Parameter	Minimum	Maximum	Notes
Static Compressive Load	—	600 N [135 lbf]	1, 2, 3
Dynamic Compressive Load	—	712 N [160 lbf]	1, 3, 4
Notes: 1. These specifications apply to uniform compressive loading in a direction normal to the processor, IHS. 2. This is the maximum static force that can be applied by the heatsink and retention solution to maintain the heatsink and processor interface. 3. These specifications are based on limited testing for design characterization. Loading limits are for the package only and do not include the limits of the processor socket. 4. Dynamic loading is defined as an 50g shock load, 2X Dynamic Acceleration Factor with a 500g maximum thermal solution.			

8.3 Package Handling Guidelines

The following table includes a list of guidelines on package handling in terms of recommended maximum loading on the processor IHS relative to a fixed substrate. These package handling loads may be experienced during heatsink removal.

Table 8-1. Package Handling Guidelines

Parameter	Maximum Recommended	Notes
Shear	311 N [70 lbf]	1, 4
Tensile	111 N [25 lbf]	2, 4
Torque	3.95 N-m [35 lbf-in]	3, 4
Notes: 1. A shear load is defined as a load applied to the IHS in a direction parallel to the IHS top surface. 2. A tensile load is defined as a pulling load applied to the IHS in a direction normal to the IHS surface. 3. A torque load is defined as a twisting load applied to the IHS in an axis of rotation normal to the IHS top surface. 4. These guidelines are based on limited testing for design characterization.		

8.4 Package Insertion Specifications

The processor can be inserted into and removed from an LGA1150 socket 15 times. The socket should meet the LGA1150 socket requirements detailed in the *LGA1150 Socket Application Guide*.

8.5 Processor Mass Specification

The typical mass of the processor is 27.0 g (0.95 oz). This mass [weight] includes all the components that are included in the package.

8.6 Processor Materials

The following table lists some of the package components and associated materials.




Table 8-2. Processor Materials

Component	Material
Integrated Heat Spreader (IHS)	Nickel Plated Copper
Substrate	Fiber Reinforced Resin
Substrate Lands	Gold Plated Copper

8.7 Processor Markings

The following figure shows the top-side markings on the processor. This diagram aids in the identification of the processor.


Figure 8-3. Processor Top-Side Markings



8.8 Processor Land Coordinates

The following figure shows the bottom view of the processor package.

Figure 8-4. Processor Package Land Coordinates



8.9 Processor Storage Specifications

The following table includes a list of the specifications for device storage in terms of maximum and minimum temperatures and relative humidity. These conditions should not be exceeded in storage or transportation.




Table 8-3. Processor Storage Specifications

Parameter	Description	Minimum	Maximum	Notes
T _{absolute storage}	The non-operating device storage temperature. Damage (latent or otherwise) may occur when subjected to for any length of time.	-55 °C	125 °C	1, 2, 3
T _{sustained storage}	The ambient storage temperature limit (in shipping media) for a sustained period of time.	-5 °C	40 °C	4, 5
RH _{sustained storage}	The maximum device storage relative humidity for a sustained period of time.	60% @ 24 °C		5, 6
TIME _{sustained storage}	A prolonged or extended period of time; typically associated with customer shelf life.	0 Months	6 Months	6
<p>Notes:</p> <ol style="list-style-type: none"> 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. Specified temperatures are based on data collected. Exceptions for surface mount reflow are specified in by applicable JEDEC standard and MAS document. Non-adherence may affect processor reliability. T_{ABSOLUTE storage} applies to the unassembled component only and does not apply to the shipping media, moisture barrier bags, or desiccant. Intel branded board products are certified to meet the following temperature and humidity limits that are given as an example only (Non-Operating Temperature Limit: -40 °C to 70 °C, Humidity: 50% to 90%, non-condensing with a maximum wet bulb of 28 °C). Post board attach storage temperature limits are not specified for non-Intel branded boards. The JEDEC, J-JSTD-020 moisture level rating and associated handling practices apply to all moisture sensitive devices removed from the moisture barrier bag. Nominal temperature and humidity conditions and durations are given and tested within the constraints imposed by T_{sustained storage} and customer shelf life in applicable intel box and bags. 				

§





9 Processor Ball and Signal Information

This chapter provides processor ball information. The following table provides the ball list by signal name.

Note: References to SA_ECC_CB[7:0] and SB_ECC_CB[7:0] are ONLY for processor SKUs that support ECC.



Table 9-1. Processor Ball List by Signal Name (Sheet 1 of 26)

Signal Name	Ball #
BCLKN	V4
BCLKP	V5
BPM#0	G39
BPM#1	J39
BPM#2	G38
BPM#3	H37
BPM#4	H38
BPM#5	J38
BPM#6	K39
BPM#7	K37
CATERR#	M36
CFG_RCOMP	H40
CFG0	AA37
CFG1	Y38
CFG10	AA34
CFG11	V37
CFG12	Y34
CFG13	U38
CFG14	W34
CFG15	V35
CFG16	Y37
CFG17	Y36
CFG18	W36
CFG19	V36
CFG2	AA36
CFG3	W38
CFG4	V39
CFG5	U39
CFG6	U40
CFG7	V38
CFG8	T40
CFG9	Y35
DBR#	G40
DDIB_TXBN0	F17
DDIB_TXBN1	G18
DDIB_TXBN2	H19
DDIB_TXBN3	G20
DDIB_TXBP0	E17
DDIB_TXBP1	F18
DDIB_TXBP2	G19
DDIB_TXBP3	F20
DDIC_TXCN0	E19
DDIC_TXCN1	D20
DDIC_TXCN2	E21
DDIC_TXCN3	D22
DDIC_TXCP0	D19

Table 9-1. Processor Ball List by Signal Name (Sheet 2 of 26)

Signal Name	Ball #
DDIC_TXCP1	C20
DDIC_TXCP2	D21
DDIC_TXCP3	C22
DDID_TXDN0	C15
DDID_TXDN1	B16
DDID_TXDN2	C17
DDID_TXDN3	B18
DDID_TXDP0	B15
DDID_TXDP1	A16
DDID_TXDP2	B17
DDID_TXDP3	A18
DISP_INT	D18
DMI_RXN0	T3
DMI_RXN1	V1
DMI_RXN2	V2
DMI_RXN3	W3
DMI_RXP0	U3
DMI_RXP1	U1
DMI_RXP2	W2
DMI_RXP3	Y3
DMI_TXN0	AA5
DMI_TXN1	AB4
DMI_TXN2	AC4
DMI_TXN3	AC2
DMI_TXP0	AA4
DMI_TXP1	AB3
DMI_TXP2	AC5
DMI_TXP3	AC1
DP_RCOMP	R4
DPLL_REF_CLKN	W6
DPLL_REF_CLKP	W5
EDP_DISP_UTIL	E16
VCCST	K9
VCCST_PWRGD	Y7
FDI_CSYNC	D16
FDIO_TX0N0	B14
FDIO_TX0N1	C13
FDIO_TX0P0	A14
FDIO_TX0P1	B13
IST_TRIGGER	C39
IVR_ERROR	R36
PECI	N37
PEG_RCOMP	P3
PEG_RXN0	F15
PEG_RXN1	E14
PEG_RXN10	F6

Table 9-1. Processor Ball List by Signal Name (Sheet 3 of 26)

Signal Name	Ball #
PEG_RXN11	G5
PEG_RXN12	H6
PEG_RXN13	J5
PEG_RXN14	K6
PEG_RXN15	L5
PEG_RXN2	F13
PEG_RXN3	E12
PEG_RXN4	F11
PEG_RXN5	G10
PEG_RXN6	F9
PEG_RXN7	G8
PEG_RXN8	D4
PEG_RXN9	E5
PEG_RXP0	E15
PEG_RXP1	D14
PEG_RXP10	F5
PEG_RXP11	G4
PEG_RXP12	H5
PEG_RXP13	J4
PEG_RXP14	K5
PEG_RXP15	L4
PEG_RXP2	E13
PEG_RXP3	D12
PEG_RXP4	E11
PEG_RXP5	F10
PEG_RXP6	E9
PEG_RXP7	F8
PEG_RXP8	D3
PEG_RXP9	E4
PEG_TXN0	B12
PEG_TXN1	C11
PEG_TXN10	G2
PEG_TXN11	H3
PEG_TXN12	J2
PEG_TXN13	K3
PEG_TXN14	M3
PEG_TXN15	L2
PEG_TXN2	D10
PEG_TXN3	C9
PEG_TXN4	D8
PEG_TXN5	C7
PEG_TXN6	B6
PEG_TXN7	C5
PEG_TXN8	E2
PEG_TXN9	F3
PEG_TXP0	A12



Table 9-1. Processor Ball List by Signal Name (Sheet 4 of 26)

Signal Name	Ball #
PEG_TXP1	B11
PEG_TXP10	G1
PEG_TXP11	H2
PEG_TXP12	J1
PEG_TXP13	K2
PEG_TXP14	M2
PEG_TXP15	L1
PEG_TXP2	C10
PEG_TXP3	B9
PEG_TXP4	C8
PEG_TXP5	B7
PEG_TXP6	A6
PEG_TXP7	B5
PEG_TXP8	E1
PEG_TXP9	F2
PM_SYNC	P36
PRDY#	L39
PREQ#	L37
PROCHOT#	K38
PWR_DEBUG	N40
PWRGOOD	AB35
RESET#	M39
RSVD	AB33
RSVD	AB36
RSVD	AB8
RSVD	AC8
RSVD	AK20
RSVD	AL20
RSVD	AT40
RSVD	AU1
RSVD	AU27
RSVD	AU39
RSVD	AV2
RSVD	AV20
RSVD	AV24
RSVD	AV29
RSVD	AW12
RSVD	AW23
RSVD	AW24
RSVD	AW27
RSVD	AY18
RSVD	H12
RSVD	H14
RSVD	H15
RSVD	J15
RSVD	J17

Table 9-1. Processor Ball List by Signal Name (Sheet 5 of 26)

Signal Name	Ball #
RSVD	J40
RSVD	J9
RSVD	L10
RSVD	L12
RSVD	M10
RSVD	M11
RSVD	M38
RSVD	N35
RSVD	P33
RSVD	R33
RSVD	R34
RSVD	T34
RSVD	T35
RSVD	T8
RSVD	U8
RSVD	W8
RSVD	Y8
RSVD_TP	A4
RSVD_TP	AV1
RSVD_TP	AW2
RSVD_TP	B3
RSVD_TP	C2
RSVD_TP	D1
RSVD_TP	H16
RSVD_TP	J10
RSVD_TP	J12
RSVD_TP	J13
RSVD_TP	J16
RSVD_TP	J8
RSVD_TP	K11
RSVD_TP	K12
RSVD_TP	K13
RSVD_TP	K8
RSVD_TP	N36
RSVD_TP	N38
RSVD_TP	P37
SA_BS0	AV12
SA_BS1	AY11
SA_BS2	AT21
SA_CAS#	AU9
SA_CK0	AY15
SA_CK1	AW15
SA_CK2	AV14
SA_CK3	AW13
SA_CKE0	AV22
SA_CKE1	AT23

Table 9-1. Processor Ball List by Signal Name (Sheet 6 of 26)

Signal Name	Ball #
SA_CKE2	AU22
SA_CKE3	AU23
SA_CKN0	AY16
SA_CKN1	AV15
SA_CKN2	AW14
SA_CKN3	AY13
SA_CS#0	AU14
SA_CS#1	AV9
SA_CS#2	AU10
SA_CS#3	AW8
SA_DIMM_VREFDQ	AB39
SA_DQ0	AD38
SA_DQ1	AD39
SA_DQ10	AK38
SA_DQ11	AK39
SA_DQ12	AH37
SA_DQ13	AH38
SA_DQ14	AK37
SA_DQ15	AK40
SA_DQ16	AM40
SA_DQ17	AM39
SA_DQ18	AP38
SA_DQ19	AP39
SA_DQ2	AF38
SA_DQ20	AM37
SA_DQ21	AM38
SA_DQ22	AP37
SA_DQ23	AP40
SA_DQ24	AV37
SA_DQ25	AW37
SA_DQ26	AU35
SA_DQ27	AV35
SA_DQ28	AT37
SA_DQ29	AU37
SA_DQ3	AF39
SA_DQ30	AT35
SA_DQ31	AW35
SA_DQ32	AY6
SA_DQ33	AU6
SA_DQ34	AV4
SA_DQ35	AU4
SA_DQ36	AW6
SA_DQ37	AV6
SA_DQ38	AW4
SA_DQ39	AY4
SA_DQ4	AD37



Table 9-1. Processor Ball List by Signal Name (Sheet 7 of 26)

Signal Name	Ball #
SA_DQ40	AR1
SA_DQ41	AR4
SA_DQ42	AN3
SA_DQ43	AN4
SA_DQ44	AR2
SA_DQ45	AR3
SA_DQ46	AN2
SA_DQ47	AN1
SA_DQ48	AL1
SA_DQ49	AL4
SA_DQ5	AD40
SA_DQ50	AJ3
SA_DQ51	AJ4
SA_DQ52	AL2
SA_DQ53	AL3
SA_DQ54	AJ2
SA_DQ55	AJ1
SA_DQ56	AG1
SA_DQ57	AG4
SA_DQ58	AE3
SA_DQ59	AE4
SA_DQ6	AF37
SA_DQ60	AG2
SA_DQ61	AG3
SA_DQ62	AE2
SA_DQ63	AE1
SA_DQ7	AF40
SA_DQ8	AH40
SA_DQ9	AH39
SA_DQSN0	AE38
SA_DQSN1	AJ38
SA_DQSN2	AN38
SA_DQSN3	AU36
SA_DQSN4	AW5
SA_DQSN5	AP2
SA_DQSN6	AK2
SA_DQSN7	AF2
SA_DQSN8	AU32
SA_DQSP0	AE39
SA_DQSP1	AJ39
SA_DQSP2	AN39
SA_DQSP3	AV36
SA_DQSP4	AV5
SA_DQSP5	AP3
SA_DQSP6	AK3
SA_DQSP7	AF3

Table 9-1. Processor Ball List by Signal Name (Sheet 8 of 26)

Signal Name	Ball #
SA_DQSP8	AV32
SA_ECC_CB0	AW33
SA_ECC_CB1	AV33
SA_ECC_CB2	AU31
SA_ECC_CB3	AV31
SA_ECC_CB4	AT33
SA_ECC_CB5	AU33
SA_ECC_CB6	AT31
SA_ECC_CB7	AW31
SA_MA0	AU13
SA_MA1	AV16
SA_MA10	AW11
SA_MA11	AV19
SA_MA12	AU19
SA_MA13	AY10
SA_MA14	AT20
SA_MA15	AU21
SA_MA2	AU16
SA_MA3	AW17
SA_MA4	AU17
SA_MA5	AW18
SA_MA6	AV17
SA_MA7	AT18
SA_MA8	AU18
SA_MA9	AT19
SA_ODT0	AW10
SA_ODT1	AY8
SA_ODT2	AW9
SA_ODT3	AU8
SA_RAS#	AU12
SA_WE#	AU11
SB_BS0	AK17
SB_BS1	AL18
SB_BS2	AW28
SB_CAS#	AP16
SB_CK0	AM20
SB_CK1	AP22
SB_CK2	AN20
SB_CK3	AP19
SB_CKE0	AW29
SB_CKE1	AY29
SB_CKE2	AU28
SB_CKE3	AU29
SB_CKN0	AM21
SB_CKN1	AP21
SB_CKN2	AN21

Table 9-1. Processor Ball List by Signal Name (Sheet 9 of 26)

Signal Name	Ball #
SB_CKN3	AP20
SB_CS#0	AP17
SB_CS#1	AN15
SB_CS#2	AN17
SB_CS#3	AL15
SB_DIMM_VREFDQ	AB40
SB_DQ0	AE34
SB_DQ1	AE35
SB_DQ10	AK31
SB_DQ11	AL31
SB_DQ12	AK34
SB_DQ13	AK35
SB_DQ14	AK32
SB_DQ15	AL32
SB_DQ16	AN34
SB_DQ17	AP34
SB_DQ18	AN31
SB_DQ19	AP31
SB_DQ2	AG35
SB_DQ20	AN35
SB_DQ21	AP35
SB_DQ22	AN32
SB_DQ23	AP32
SB_DQ24	AM29
SB_DQ25	AM28
SB_DQ26	AR29
SB_DQ27	AR28
SB_DQ28	AL29
SB_DQ29	AL28
SB_DQ3	AH35
SB_DQ30	AP29
SB_DQ31	AP28
SB_DQ32	AR12
SB_DQ33	AP12
SB_DQ34	AL13
SB_DQ35	AL12
SB_DQ36	AR13
SB_DQ37	AP13
SB_DQ38	AM13
SB_DQ39	AM12
SB_DQ4	AD34
SB_DQ40	AR9
SB_DQ41	AP9
SB_DQ42	AR6
SB_DQ43	AP6
SB_DQ44	AR10



Table 9-1. Processor Ball List by Signal Name (Sheet 10 of 26)

Signal Name	Ball #
SB_DQ45	AP10
SB_DQ46	AR7
SB_DQ47	AP7
SB_DQ48	AM9
SB_DQ49	AL9
SB_DQ5	AD35
SB_DQ50	AL6
SB_DQ51	AL7
SB_DQ52	AM10
SB_DQ53	AL10
SB_DQ54	AM6
SB_DQ55	AM7
SB_DQ56	AH6
SB_DQ57	AH7
SB_DQ58	AE6
SB_DQ59	AE7
SB_DQ6	AG34
SB_DQ60	AJ6
SB_DQ61	AJ7
SB_DQ62	AF6
SB_DQ63	AF7
SB_DQ7	AH34
SB_DQ8	AL34
SB_DQ9	AL35
SB_DQS0	AF35
SB_DQS1	AL33
SB_DQS2	AP33
SB_DQS3	AN28
SB_DQS4	AN12
SB_DQS5	AP8
SB_DQS6	AL8
SB_DQS7	AG7
SB_DQS8	AN25
SB_DQSN0	AF34
SB_DQSN1	AK33
SB_DQSN2	AN33
SB_DQSN3	AN29
SB_DQSN4	AN13
SB_DQSN5	AR8
SB_DQSN6	AM8
SB_DQSN7	AG6
SB_DQSN8	AN26
SB_ECC_CB0	AM26
SB_ECC_CB1	AM25
SB_ECC_CB2	AP25
SB_ECC_CB3	AP26

Table 9-1. Processor Ball List by Signal Name (Sheet 11 of 26)

Signal Name	Ball #
SB_ECC_CB4	AL26
SB_ECC_CB5	AL25
SB_ECC_CB6	AR26
SB_ECC_CB7	AR25
SB_MA0	AL19
SB_MA1	AK23
SB_MA10	AP18
SB_MA11	AY25
SB_MA12	AV26
SB_MA13	AR15
SB_MA14	AV27
SB_MA15	AY28
SB_MA2	AM22
SB_MA3	AM23
SB_MA4	AP23
SB_MA5	AL23
SB_MA6	AY24
SB_MA7	AV25
SB_MA8	AU26
SB_MA9	AW25
SB_ODT0	AM17
SB_ODT1	AL16
SB_ODT2	AM16
SB_ODT3	AK15
SB_RAS#	AM18
SB_WE#	AK16
SKTOCC#	D38
SM_DRAMPWROK	AK21
SM_DRAMRST#	AK22
SM_RCOMP0	R1
SM_RCOMP1	P1
SM_RCOMP2	R2
SM_VREF	AB38
SSC_DPLL_REF_CLK N	U5
SSC_DPLL_REF_CLKP	U6
TCK	D39
TDI	F38
TDO	F39
TESTLO_N5	N5
TESTLO_P6	P6
THERMTRIP#	F37
TMS	E39
TRST#	E37
VCC	A24
VCC	A25

Table 9-1. Processor Ball List by Signal Name (Sheet 12 of 26)

Signal Name	Ball #
VCC	A26
VCC	A27
VCC	A28
VCC	A29
VCC	A30
VCC	B25
VCC	B27
VCC	B29
VCC	B31
VCC	B33
VCC	B35
VCC	C24
VCC	C25
VCC	C26
VCC	C27
VCC	C28
VCC	C29
VCC	C30
VCC	C31
VCC	C32
VCC	C33
VCC	C34
VCC	C35
VCC	D25
VCC	D27
VCC	D29
VCC	D31
VCC	D33
VCC	D35
VCC	E24
VCC	E25
VCC	E26
VCC	E27
VCC	E28
VCC	E29
VCC	E30
VCC	E31
VCC	E32
VCC	E33
VCC	E34
VCC	E35
VCC	F23
VCC	F25
VCC	F27
VCC	F29
VCC	F31



Table 9-1. Processor Ball List by Signal Name (Sheet 13 of 26)

Signal Name	Ball #
VCC	F33
VCC	F35
VCC	G22
VCC	G23
VCC	G24
VCC	G25
VCC	G26
VCC	G27
VCC	G28
VCC	G29
VCC	G30
VCC	G31
VCC	G32
VCC	G33
VCC	G34
VCC	G35
VCC	H23
VCC	H25
VCC	H27
VCC	H29
VCC	H31
VCC	H33
VCC	H35
VCC	J21
VCC	J22
VCC	J23
VCC	J24
VCC	J25
VCC	J26
VCC	J27
VCC	J28
VCC	J29
VCC	J30
VCC	J31
VCC	J32
VCC	J33
VCC	J34
VCC	J35
VCC	K19
VCC	K21
VCC	K23
VCC	K25
VCC	K27
VCC	K29
VCC	K31
VCC	K33

Table 9-1. Processor Ball List by Signal Name (Sheet 14 of 26)

Signal Name	Ball #
VCC	K35
VCC	L15
VCC	L16
VCC	L17
VCC	L18
VCC	L19
VCC	L20
VCC	L21
VCC	L22
VCC	L23
VCC	L24
VCC	L25
VCC	L26
VCC	L27
VCC	L28
VCC	L29
VCC	L30
VCC	L31
VCC	L32
VCC	L33
VCC	L34
VCC	M13
VCC	M15
VCC	M17
VCC	M19
VCC	M21
VCC	M23
VCC	M25
VCC	M27
VCC	M29
VCC	M33
VCC	M8
VCC	P8
VCC_SENSE	E40
VCCIO_OUT	L40
VCOMP_OUT	P4
VDDQ	AJ12
VDDQ	AJ13
VDDQ	AJ15
VDDQ	AJ17
VDDQ	AJ20
VDDQ	AJ21
VDDQ	AJ24
VDDQ	AJ25
VDDQ	AJ28
VDDQ	AJ29

Table 9-1. Processor Ball List by Signal Name (Sheet 15 of 26)

Signal Name	Ball #
VDDQ	AJ9
VDDQ	AT17
VDDQ	AT22
VDDQ	AU15
VDDQ	AU20
VDDQ	AU24
VDDQ	AV10
VDDQ	AV11
VDDQ	AV13
VDDQ	AV18
VDDQ	AV23
VDDQ	AV8
VDDQ	AW16
VDDQ	AY12
VDDQ	AY14
VDDQ	AY9
VIDALERT#	B37
VIDSCLK	C38
VIDSOUT	C37
VSS	A11
VSS	A13
VSS	A15
VSS	A17
VSS	A23
VSS	A5
VSS	A7
VSS	AA3
VSS	AA33
VSS	AA35
VSS	AA38
VSS	AA6
VSS	AA7
VSS	AA8
VSS	AB34
VSS	AB37
VSS	AB5
VSS	AB6
VSS	AB7
VSS	AC3
VSS	AC33
VSS	AC34
VSS	AC35
VSS	AC36
VSS	AC37
VSS	AC38
VSS	AC39



Table 9-1. Processor Ball List by Signal Name (Sheet 16 of 26)

Signal Name	Ball #
VSS	AC40
VSS	AC6
VSS	AC7
VSS	AD1
VSS	AD2
VSS	AD3
VSS	AD33
VSS	AD36
VSS	AD4
VSS	AD5
VSS	AD6
VSS	AD7
VSS	AD8
VSS	AE33
VSS	AE36
VSS	AE37
VSS	AE40
VSS	AE5
VSS	AE8
VSS	AF1
VSS	AF33
VSS	AF36
VSS	AF4
VSS	AF5
VSS	AF8
VSS	AG33
VSS	AG36
VSS	AG37
VSS	AG38
VSS	AG39
VSS	AG40
VSS	AG5
VSS	AG8
VSS	AH1
VSS	AH2
VSS	AH3
VSS	AH33
VSS	AH36
VSS	AH4
VSS	AH5
VSS	AH8
VSS	AJ11
VSS	AJ14
VSS	AJ16
VSS	AJ18
VSS	AJ19

Table 9-1. Processor Ball List by Signal Name (Sheet 17 of 26)

Signal Name	Ball #
VSS	AJ22
VSS	AJ23
VSS	AJ26
VSS	AJ27
VSS	AJ30
VSS	AJ31
VSS	AJ32
VSS	AJ33
VSS	AJ34
VSS	AJ35
VSS	AJ36
VSS	AJ37
VSS	AJ40
VSS	AJ5
VSS	AJ8
VSS	AK1
VSS	AK10
VSS	AK11
VSS	AK12
VSS	AK13
VSS	AK14
VSS	AK18
VSS	AK19
VSS	AK24
VSS	AK25
VSS	AK26
VSS	AK27
VSS	AK28
VSS	AK29
VSS	AK30
VSS	AK36
VSS	AK4
VSS	AK5
VSS	AK6
VSS	AK7
VSS	AK8
VSS	AK9
VSS	AL11
VSS	AL14
VSS	AL17
VSS	AL21
VSS	AL22
VSS	AL24
VSS	AL27
VSS	AL30
VSS	AL36

Table 9-1. Processor Ball List by Signal Name (Sheet 18 of 26)

Signal Name	Ball #
VSS	AL37
VSS	AL38
VSS	AL39
VSS	AL40
VSS	AL5
VSS	AM1
VSS	AM11
VSS	AM14
VSS	AM15
VSS	AM19
VSS	AM2
VSS	AM24
VSS	AM27
VSS	AM3
VSS	AM30
VSS	AM31
VSS	AM32
VSS	AM33
VSS	AM34
VSS	AM35
VSS	AM36
VSS	AM4
VSS	AM5
VSS	AN10
VSS	AN11
VSS	AN14
VSS	AN16
VSS	AN18
VSS	AN19
VSS	AN22
VSS	AN23
VSS	AN24
VSS	AN27
VSS	AN30
VSS	AN36
VSS	AN37
VSS	AN40
VSS	AN5
VSS	AN6
VSS	AN7
VSS	AN8
VSS	AN9
VSS	AP1
VSS	AP11
VSS	AP14
VSS	AP15



Table 9-1. Processor Ball List by Signal Name (Sheet 19 of 26)

Signal Name	Ball #
VSS	AP24
VSS	AP27
VSS	AP30
VSS	AP36
VSS	AP4
VSS	AP5
VSS	AR11
VSS	AR14
VSS	AR16
VSS	AR17
VSS	AR18
VSS	AR19
VSS	AR20
VSS	AR21
VSS	AR22
VSS	AR23
VSS	AR24
VSS	AR27
VSS	AR30
VSS	AR31
VSS	AR32
VSS	AR33
VSS	AR34
VSS	AR35
VSS	AR36
VSS	AR37
VSS	AR38
VSS	AR39
VSS	AR40
VSS	AR5
VSS	AT1
VSS	AT10
VSS	AT11
VSS	AT12
VSS	AT13
VSS	AT14
VSS	AT15
VSS	AT16
VSS	AT2
VSS	AT24
VSS	AT25
VSS	AT26
VSS	AT27
VSS	AT28
VSS	AT29
VSS	AT3

Table 9-1. Processor Ball List by Signal Name (Sheet 20 of 26)

Signal Name	Ball #
VSS	AT30
VSS	AT32
VSS	AT34
VSS	AT36
VSS	AT38
VSS	AT39
VSS	AT4
VSS	AT5
VSS	AT6
VSS	AT7
VSS	AT8
VSS	AT9
VSS	AU2
VSS	AU25
VSS	AU3
VSS	AU30
VSS	AU34
VSS	AU38
VSS	AU5
VSS	AU7
VSS	AV21
VSS	AV28
VSS	AV3
VSS	AV30
VSS	AV34
VSS	AV38
VSS	AV7
VSS	AW26
VSS	AW3
VSS	AW30
VSS	AW32
VSS	AW34
VSS	AW36
VSS	AW7
VSS	AY17
VSS	AY23
VSS	AY26
VSS	AY27
VSS	AY30
VSS	AY5
VSS	AY7
VSS	B10
VSS	B23
VSS	B24
VSS	B26
VSS	B28

Table 9-1. Processor Ball List by Signal Name (Sheet 21 of 26)

Signal Name	Ball #
VSS	B30
VSS	B32
VSS	B34
VSS	B36
VSS	B4
VSS	B8
VSS	C12
VSS	C14
VSS	C16
VSS	C18
VSS	C19
VSS	C21
VSS	C23
VSS	C3
VSS	C36
VSS	C4
VSS	C6
VSS	D11
VSS	D13
VSS	D15
VSS	D17
VSS	D2
VSS	D23
VSS	D24
VSS	D26
VSS	D28
VSS	D30
VSS	D32
VSS	D34
VSS	D36
VSS	D37
VSS	D5
VSS	D6
VSS	D7
VSS	D9
VSS	E10
VSS	E18
VSS	E20
VSS	E22
VSS	E23
VSS	E3
VSS	E36
VSS	E38
VSS	E6
VSS	E7
VSS	E8



Table 9-1. Processor Ball List by Signal Name (Sheet 22 of 26)

Signal Name	Ball #
VSS	F1
VSS	F12
VSS	F14
VSS	F16
VSS	F19
VSS	F21
VSS	F22
VSS	F24
VSS	F26
VSS	F28
VSS	F30
VSS	F32
VSS	F34
VSS	F36
VSS	F4
VSS	F7
VSS	G11
VSS	G12
VSS	G13
VSS	G14
VSS	G15
VSS	G16
VSS	G17
VSS	G21
VSS	G3
VSS	G36
VSS	G37
VSS	G6
VSS	G7
VSS	G9
VSS	H1
VSS	H10
VSS	H11
VSS	H13
VSS	H17
VSS	H18
VSS	H20
VSS	H21
VSS	H22
VSS	H24
VSS	H26
VSS	H28
VSS	H30
VSS	H32
VSS	H34
VSS	H36

Table 9-1. Processor Ball List by Signal Name (Sheet 23 of 26)

Signal Name	Ball #
VSS	H39
VSS	H4
VSS	H7
VSS	H8
VSS	H9
VSS	J11
VSS	J14
VSS	J18
VSS	J19
VSS	J20
VSS	J3
VSS	J36
VSS	J37
VSS	J6
VSS	J7
VSS	K1
VSS	K10
VSS	K14
VSS	K15
VSS	K16
VSS	K17
VSS	K18
VSS	K20
VSS	K22
VSS	K24
VSS	K26
VSS	K28
VSS	K30
VSS	K32
VSS	K34
VSS	K36
VSS	K4
VSS	K40
VSS	K7
VSS	L11
VSS	L13
VSS	L14
VSS	L3
VSS	L35
VSS	L36
VSS	L38
VSS	L6
VSS	L7
VSS	L8
VSS	L9
VSS	M1

Table 9-1. Processor Ball List by Signal Name (Sheet 24 of 26)

Signal Name	Ball #
VSS	M12
VSS	M14
VSS	M16
VSS	M18
VSS	M20
VSS	M22
VSS	M24
VSS	M26
VSS	M28
VSS	M30
VSS	M32
VSS	M34
VSS	M35
VSS	M37
VSS	M4
VSS	M40
VSS	M5
VSS	M6
VSS	M7
VSS	M9
VSS	N1
VSS	N2
VSS	N3
VSS	N33
VSS	N34
VSS	N39
VSS	N4
VSS	N6
VSS	N7
VSS	N8
VSS	P2
VSS	P34
VSS	P35
VSS	P38
VSS	P39
VSS	P40
VSS	P5
VSS	P7
VSS	R3
VSS	R35
VSS	R37
VSS	R38
VSS	R39
VSS	R40
VSS	R5
VSS	R6




Table 9-1. Processor Ball List by Signal Name (Sheet 25 of 26)

Signal Name	Ball #
VSS	R7
VSS	R8
VSS	T1
VSS	T2
VSS	T33
VSS	T36
VSS	T37
VSS	T38
VSS	T39
VSS	T4
VSS	T5
VSS	T6
VSS	T7
VSS	U2
VSS	U33
VSS	U34
VSS	U35
VSS	U36
VSS	U37
VSS	U4
VSS	U7
VSS	V3
VSS	V33
VSS	V34
VSS	V40
VSS	V6
VSS	V7
VSS	V8
VSS	W1
VSS	W33
VSS	W35
VSS	W37
VSS	W4
VSS	W7
VSS	Y33
VSS	Y4
VSS	Y5
VSS	Y6
VSS_NCTF	AU40
VSS_NCTF	AV39
VSS_NCTF	AW38
VSS_NCTF	AY3
VSS_NCTF	B38
VSS_NCTF	B39

Table 9-1. Processor Ball List by Signal Name (Sheet 26 of 26)

Signal Name	Ball #
VSS_NCTF	C40
VSS_NCTF	D40
VSS_SENSE	F40

§



10 DDR Data Swizzling

The processors do not implement DDR data swizzling.

§

