

Intel® Xeon® Processor 3500 Series

Datasheet, Volume 2

March 2009



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® Xeon® Processor 3500 Series may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details. Over time processor numbers will increment based on changes in clock, speed, cache, FSB, or other features, and increments are not intended to represent proportional or quantitative increases in any particular feature. Current roadmap processor number progression is not necessarily representative of future roadmaps. See www.intel.com/products/processor_number for details.

Hyper-Threading Technology requires a computer system with a processor supporting HT Technology and an HT Technology-enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. For more information including details on which processors support HT Technology, see http://www.intel.com/products/ht/hyperthreading_more.htm.

Enabling Execute Disable Bit functionality requires a PC with a processor with Execute Disable Bit capability and a supporting operating system. Check with your PC manufacturer on whether your system delivers Execute Disable Bit functionality.

64-bit computing on Intel architecture requires a computer system with a processor, chipset, BIOS, operating system, device drivers and applications enabled for Intel® 64 architecture. Performance will vary depending on your hardware and software configurations. Consult with your system vendor for more information.

Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain computer system software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

Intel® Turbo Boost Technology requires a PC with a processor with Intel Turbo Boost Technology capability. Intel Turbo Boost Technology performance varies depending on hardware, software and overall system configuration. Check with your PC manufacturer on whether your system delivers Intel Turbo Boost Technology. For more information, see www.intel.com.

Enhanced Intel SpeedStep® Technology. See the <http://processorfinder.intel.com> or contact your Intel representative for more information.

Intel, Xeon, Enhanced Intel SpeedStep Technology and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

*Other brands and names are the property of their respective owners.

Copyright © 2009, Intel Corporation.



Contents


1	Introduction	11
1.1	Terminology	11
1.1.1	Processor Terminology	11
1.2	References	13
2	Register Description	15
2.1	Register Terminology	15
2.2	Platform Configuration Structure	16
2.3	Device Mapping	17
2.4	Detailed Configuration Space Maps	18
2.5	PCI Standard Registers	36
2.5.1	VID - Vendor Identification Register	36
2.5.2	DID - Device Identification Register	36
2.5.3	RID - Revision Identification Register	37
2.5.4	CCR - Class Code Register	37
2.5.5	HDR - Header Type Register	38
2.5.6	SID/SVID - Subsystem Identity/Subsystem Vendor Identification Register	38
2.5.7	PCICMD - Command Register	39
2.5.8	PCISTS - PCI Status Register	40
2.6	SAD - System Address Decoder Registers	41
2.6.1	SAD_PAM0123	41
2.6.2	SAD_PAM456	43
2.6.3	SAD_HEN	44
2.6.4	SAD_SMRAM	44
2.6.5	SAD_PCIEBAR	45
2.6.6	SAD_DRAM_RULE_0, SAD_DRAM_RULE_1, SAD_DRAM_RULE_2, SAD_DRAM_RULE_3, SAD_DRAM_RULE_4, SAD_DRAM_RULE_5, SAD_DRAM_RULE_6, SAD_DRAM_RULE_7	46
2.6.7	SAD_INTERLEAVE_LIST_0, SAD_INTERLEAVE_LIST_1, SAD_INTERLEAVE_LIST_2, SAD_INTERLEAVE_LIST_3, SAD_INTERLEAVE_LIST_4, SAD_INTERLEAVE_LIST_5, SAD_INTERLEAVE_LIST_6, SAD_INTERLEAVE_LIST_7	46
2.7	Intel® QuickPath Interconnect Link Registers	47
2.7.1	QPI_QPILCL_L0, QPI_QPILCL_L1	47
2.8	Integrated Memory Controller Control Registers	48
2.8.1	MC_CONTROL	48
2.8.2	MC_STATUS	49
2.8.3	MC_SMI_SPARE_DIMM_ERROR_STATUS	50
2.8.4	MC_SMI_SPARE_CNTRL	51
2.8.5	MC_RESET_CONTROL	51
2.8.6	MC_CHANNEL_MAPPER	52
2.8.7	MC_MAX_DOD	53
2.8.8	MC_RD_CRDT_INIT	54
2.8.9	MC_CRDT_WR_THLD	55
2.8.10	MC_SCRUBADDR_LO	55
2.8.11	MC_SCRUBADDR_HI	56
2.9	TAD – Target Address Decoder Registers	57
2.9.1	TAD_DRAM_RULE_0, TAD_DRAM_RULE_1, TAD_DRAM_RULE_2, TAD_DRAM_RULE_3, TAD_DRAM_RULE_4, TAD_DRAM_RULE_5, TAD_DRAM_RULE_6, TAD_DRAM_RULE_7	57



2.9.2	TAD_INTERLEAVE_LIST_0, TAD_INTERLEAVE_LIST_1 TAD_INTERLEAVE_LIST_2, TAD_INTERLEAVE_LIST_3 TAD_INTERLEAVE_LIST_4, TAD_INTERLEAVE_LIST_5 TAD_INTERLEAVE_LIST_6, TAD_INTERLEAVE_LIST_7.....	58
2.10	Integrated Memory Controller Channel Control Registers.....	59
2.10.1	MC_CHANNEL_0_DIMM_RESET_CMD MC_CHANNEL_1_DIMM_RESET_CMD MC_CHANNEL_2_DIMM_RESET_CMD	59
2.10.2	MC_CHANNEL_0_DIMM_INIT_CMD MC_CHANNEL_1_DIMM_INIT_CMD MC_CHANNEL_2_DIMM_INIT_CMD	60
2.10.3	MC_CHANNEL_0_DIMM_INIT_PARAMS MC_CHANNEL_1_DIMM_INIT_PARAMS MC_CHANNEL_2_DIMM_INIT_PARAMS	61
2.10.4	MC_CHANNEL_0_DIMM_INIT_STATUS MC_CHANNEL_1_DIMM_INIT_STATUS MC_CHANNEL_2_DIMM_INIT_STATUS	62
2.10.5	MC_CHANNEL_0_DDR3CMD MC_CHANNEL_1_DDR3CMD MC_CHANNEL_2_DDR3CMD.....	63
2.10.6	MC_CHANNEL_0_REFRESH_THROTTLE_SUPPORT MC_CHANNEL_1_REFRESH_THROTTLE_SUPPORT MC_CHANNEL_2_REFRESH_THROTTLE_SUPPORT	64
2.10.7	MC_CHANNEL_0_MRS_VALUE_0_1 MC_CHANNEL_1_MRS_VALUE_0_1 MC_CHANNEL_2_MRS_VALUE_0_1	64
2.10.8	MC_CHANNEL_0_MRS_VALUE_2 MC_CHANNEL_1_MRS_VALUE_2 MC_CHANNEL_2_MRS_VALUE_2	65
2.10.9	MC_CHANNEL_0_RANK_PRESENT MC_CHANNEL_1_RANK_PRESENT MC_CHANNEL_2_RANK_PRESENT.....	65
2.10.10	MC_CHANNEL_0_RANK_TIMING_A MC_CHANNEL_1_RANK_TIMING_A MC_CHANNEL_2_RANK_TIMING_A	66
2.10.11	MC_CHANNEL_0_RANK_TIMING_B MC_CHANNEL_1_RANK_TIMING_B MC_CHANNEL_2_RANK_TIMING_B	69
2.10.12	MC_CHANNEL_0_BANK_TIMING MC_CHANNEL_1_BANK_TIMING MC_CHANNEL_2_BANK_TIMING.....	70
2.10.13	MC_CHANNEL_0_REFRESH_TIMING MC_CHANNEL_1_REFRESH_TIMING MC_CHANNEL_2_REFRESH_TIMING	70
2.10.14	MC_CHANNEL_0_CKE_TIMING MC_CHANNEL_1_CKE_TIMING MC_CHANNEL_2_CKE_TIMING	71
2.10.15	MC_CHANNEL_0_ZQ_TIMING MC_CHANNEL_1_ZQ_TIMING MC_CHANNEL_2_ZQ_TIMING	71
2.10.16	MC_CHANNEL_0_RCOMP_PARAMS MC_CHANNEL_1_RCOMP_PARAMS MC_CHANNEL_2_RCOMP_PARAMS.....	72
2.10.17	MC_CHANNEL_0_ODT_PARAMS1 MC_CHANNEL_1_ODT_PARAMS1 MC_CHANNEL_2_ODT_PARAMS1	72
2.10.18	MC_CHANNEL_0_ODT_PARAMS2 MC_CHANNEL_1_ODT_PARAMS2 MC_CHANNEL_2_ODT_PARAMS2	73



2.10.19	MC_CHANNEL_0_ODT_MATRIX_RANK_0_3_RD	
	MC_CHANNEL_1_ODT_MATRIX_RANK_0_3_RD	
	MC_CHANNEL_2_ODT_MATRIX_RANK_0_3_RD.....	73
2.10.20	MC_CHANNEL_0_ODT_MATRIX_RANK_4_7_RD	
	MC_CHANNEL_1_ODT_MATRIX_RANK_4_7_RD	
	MC_CHANNEL_2_ODT_MATRIX_RANK_4_7_RD.....	74
2.10.21	MC_CHANNEL_0_ODT_MATRIX_RANK_0_3_WR	
	MC_CHANNEL_1_ODT_MATRIX_RANK_0_3_WR	
	MC_CHANNEL_2_ODT_MATRIX_RANK_0_3_WR.....	74
2.10.22	MC_CHANNEL_0_ODT_MATRIX_RANK_4_7_WR	
	MC_CHANNEL_1_ODT_MATRIX_RANK_4_7_WR	
	MC_CHANNEL_2_ODT_MATRIX_RANK_4_7_WR.....	74
2.10.23	MC_CHANNEL_0_WAQ_PARAMS	
	MC_CHANNEL_1_WAQ_PARAMS	
	MC_CHANNEL_2_WAQ_PARAMS	75
2.10.24	MC_CHANNEL_0_SCHEDULER_PARAMS	
	MC_CHANNEL_1_SCHEDULER_PARAMS	
	MC_CHANNEL_2_SCHEDULER_PARAMS	76
2.10.25	MC_CHANNEL_0_MAINTENANCE_OPS	
	MC_CHANNEL_1_MAINTENANCE_OPS	
	MC_CHANNEL_2_MAINTENANCE_OPS	76
2.10.26	MC_CHANNEL_0_TX_BG_SETTINGS	
	MC_CHANNEL_1_TX_BG_SETTINGS	
	MC_CHANNEL_2_TX_BG_SETTINGS.....	77
2.10.27	MC_CHANNEL_0_RX_BGF_SETTINGS	
	MC_CHANNEL_1_RX_BGF_SETTINGS	
	MC_CHANNEL_2_RX_BGF_SETTINGS	77
2.10.28	MC_CHANNEL_0_EW_BGF_SETTINGS	
	MC_CHANNEL_1_EW_BGF_SETTINGS	
	MC_CHANNEL_2_EW_BGF_SETTINGS.....	78
2.10.29	MC_CHANNEL_0_EW_BGF_OFFSET_SETTINGS	
	MC_CHANNEL_1_EW_BGF_OFFSET_SETTINGS	
	MC_CHANNEL_2_EW_BGF_OFFSET_SETTINGS	78
2.10.30	MC_CHANNEL_0_ROUND_TRIP_LATENCY	
	MC_CHANNEL_1_ROUND_TRIP_LATENCY	
	MC_CHANNEL_2_ROUND_TRIP_LATENCY.....	78
2.10.31	MC_CHANNEL_0_PAGETABLE_PARAMS1	
	MC_CHANNEL_1_PAGETABLE_PARAMS1	
	MC_CHANNEL_2_PAGETABLE_PARAMS1	79
2.10.32	MC_CHANNEL_0_PAGETABLE_PARAMS2	
	MC_CHANNEL_1_PAGETABLE_PARAMS2	
	MC_CHANNEL_2_PAGETABLE_PARAMS2	79
2.10.33	MC_TX_BG_CMD_DATA_RATIO_SETTINGS_CH0	
	MC_TX_BG_CMD_DATA_RATIO_SETTINGS_CH1	
	MC_TX_BG_CMD_DATA_RATIO_SETTINGS_CH2	80
2.10.34	MC_TX_BG_CMD_OFFSET_SETTINGS_CH0	
	MC_TX_BG_CMD_OFFSET_SETTINGS_CH1	
	MC_TX_BG_CMD_OFFSET_SETTINGS_CH2.....	80
2.10.35	MC_TX_BG_DATA_OFFSET_SETTINGS_CH0	
	MC_TX_BG_DATA_OFFSET_SETTINGS_CH1	
	MC_TX_BG_DATA_OFFSET_SETTINGS_CH2	80
2.10.36	MC_CHANNEL_0_ADDR_MATCH	
	MC_CHANNEL_1_ADDR_MATCH	
	MC_CHANNEL_2_ADDR_MATCH.....	81
2.10.37	MC_CHANNEL_0_ECC_ERROR_MASK	
	MC_CHANNEL_1_ECC_ERROR_MASK	
	MC_CHANNEL_2_ECC_ERROR_MASK.....	82
2.10.38	MC_CHANNEL_0_ECC_ERROR_INJECT	
	MC_CHANNEL_1_ECC_ERROR_INJECT	
	MC_CHANNEL_2_ECC_ERROR_INJECT.....	82



2.10.39	Error Injection Implementation	83
2.11	Integrated Memory Controller Channel Address Registers.....	84
2.11.1	MC_DOD_CH0_0, MC_DOD_CH0_1, MC_DOD_CH0_2	84
2.11.2	MC_DOD_CH1_0, MC_DOD_CH1_1, MC_DOD_CH1_2	85
2.11.3	MC_DOD_CH2_0, MC_DOD_CH2_1, MC_DOD_CH2_2	86
2.11.4	MC_SAG_CH0_0, MC_SAG_CH0_1, MC_SAG_CH0_2 MC_SAG_CH0_3, MC_SAG_CH0_4, MC_SAG_CH0_5 MC_SAG_CH0_6, MC_SAG_CH0_7, MC_SAG_CH1_0 MC_SAG_CH1_1, MC_SAG_CH1_2, MC_SAG_CH1_3 MC_SAG_CH1_4, MC_SAG_CH1_5, MC_SAG_CH1_6 MC_SAG_CH1_7, MC_SAG_CH2_0, MC_SAG_CH2_1 MC_SAG_CH2_2, MC_SAG_CH2_3, MC_SAG_CH2_4 MC_SAG_CH2_5, MC_SAG_CH2_6, MC_SAG_CH2_7.....	87
2.12	Integrated Memory Controller Channel Rank Registers.....	88
2.12.1	MC_RIR_LIMIT_CH0_0, MC_RIR_LIMIT_CH0_1 MC_RIR_LIMIT_CH0_2, MC_RIR_LIMIT_CH0_3 MC_RIR_LIMIT_CH0_4, MC_RIR_LIMIT_CH0_5 MC_RIR_LIMIT_CH0_6, MC_RIR_LIMIT_CH0_7 MC_RIR_LIMIT_CH1_0, MC_RIR_LIMIT_CH1_1 MC_RIR_LIMIT_CH1_2, MC_RIR_LIMIT_CH1_3 MC_RIR_LIMIT_CH1_4, MC_RIR_LIMIT_CH1_5 MC_RIR_LIMIT_CH1_6, MC_RIR_LIMIT_CH1_7 MC_RIR_LIMIT_CH2_0, MC_RIR_LIMIT_CH2_1 MC_RIR_LIMIT_CH2_2, MC_RIR_LIMIT_CH2_3 MC_RIR_LIMIT_CH2_4, MC_RIR_LIMIT_CH2_5 MC_RIR_LIMIT_CH2_6, MC_RIR_LIMIT_CH2_7	88
2.12.2	MC_RIR_WAY_CH0_0, MC_RIR_WAY_CH0_1 MC_RIR_WAY_CH0_2, MC_RIR_WAY_CH0_3 MC_RIR_WAY_CH0_4, MC_RIR_WAY_CH0_5 MC_RIR_WAY_CH0_6, MC_RIR_WAY_CH0_7 MC_RIR_WAY_CH0_8, MC_RIR_WAY_CH0_9 MC_RIR_WAY_CH0_10, MC_RIR_WAY_CH0_11 MC_RIR_WAY_CH0_12, MC_RIR_WAY_CH0_13 MC_RIR_WAY_CH0_14, MC_RIR_WAY_CH0_15 MC_RIR_WAY_CH0_16, MC_RIR_WAY_CH0_17 MC_RIR_WAY_CH0_18, MC_RIR_WAY_CH0_19 MC_RIR_WAY_CH0_20, MC_RIR_WAY_CH0_21 MC_RIR_WAY_CH0_22, MC_RIR_WAY_CH0_23 MC_RIR_WAY_CH0_24, MC_RIR_WAY_CH0_25 MC_RIR_WAY_CH0_26, MC_RIR_WAY_CH0_27 MC_RIR_WAY_CH0_28, MC_RIR_WAY_CH0_29 MC_RIR_WAY_CH0_30, MC_RIR_WAY_CH0_31	89
2.12.3	MC_RIR_WAY_CH1_0, MC_RIR_WAY_CH1_1 MC_RIR_WAY_CH1_2, MC_RIR_WAY_CH1_3 MC_RIR_WAY_CH1_4, MC_RIR_WAY_CH1_5 MC_RIR_WAY_CH1_6, MC_RIR_WAY_CH1_7 MC_RIR_WAY_CH1_8, MC_RIR_WAY_CH1_9 MC_RIR_WAY_CH1_10, MC_RIR_WAY_CH1_11 MC_RIR_WAY_CH1_12, MC_RIR_WAY_CH1_13 MC_RIR_WAY_CH1_14, MC_RIR_WAY_CH1_15 MC_RIR_WAY_CH1_16, MC_RIR_WAY_CH1_17 MC_RIR_WAY_CH1_18, MC_RIR_WAY_CH1_19 MC_RIR_WAY_CH1_20, MC_RIR_WAY_CH1_21 MC_RIR_WAY_CH1_22, MC_RIR_WAY_CH1_23 MC_RIR_WAY_CH1_24, MC_RIR_WAY_CH1_25 MC_RIR_WAY_CH1_26, MC_RIR_WAY_CH1_27 MC_RIR_WAY_CH1_28, MC_RIR_WAY_CH1_29 MC_RIR_WAY_CH1_30, MC_RIR_WAY_CH1_31	90



2.12.4	MC_RIR_WAY_CH2_0, MC_RIR_WAY_CH2_1 MC_RIR_WAY_CH2_2, MC_RIR_WAY_CH2_3 MC_RIR_WAY_CH2_4, MC_RIR_WAY_CH2_5 MC_RIR_WAY_CH2_6, MC_RIR_WAY_CH2_7 MC_RIR_WAY_CH2_8, MC_RIR_WAY_CH2_9 MC_RIR_WAY_CH2_10, MC_RIR_WAY_CH2_11 MC_RIR_WAY_CH2_12, MC_RIR_WAY_CH2_13 MC_RIR_WAY_CH2_14, MC_RIR_WAY_CH2_15 MC_RIR_WAY_CH2_16, MC_RIR_WAY_CH2_17 MC_RIR_WAY_CH2_18, MC_RIR_WAY_CH2_19 MC_RIR_WAY_CH2_20, MC_RIR_WAY_CH2_21 MC_RIR_WAY_CH2_22, MC_RIR_WAY_CH2_23 MC_RIR_WAY_CH2_24, MC_RIR_WAY_CH2_25 MC_RIR_WAY_CH2_26, MC_RIR_WAY_CH2_27 MC_RIR_WAY_CH2_28, MC_RIR_WAY_CH2_29 MC_RIR_WAY_CH2_30, MC_RIR_WAY_CH2_31	91
2.13	Memory Thermal Control	92
2.13.1	MC_THERMAL_CONTROLO MC_THERMAL_CONTROL1 MC_THERMAL_CONTROL2	92
2.13.2	MC_THERMAL_STATUS0 MC_THERMAL_STATUS1 MC_THERMAL_STATUS2	92
2.13.3	MC_THERMAL_DEFEATURE0 MC_THERMAL_DEFEATURE1 MC_THERMAL_DEFEATURE2	93
2.13.4	MC_THERMAL_PARAMS_A0 MC_THERMAL_PARAMS_A1 MC_THERMAL_PARAMS_A2	93
2.13.5	MC_THERMAL_PARAMS_B0 MC_THERMAL_PARAMS_B1 MC_THERMAL_PARAMS_B2	94
2.13.6	MC_COOLING_COEFO MC_COOLING_COEF1 MC_COOLING_COEF2	94
2.13.7	MC_CLOSED_LOOP0 MC_CLOSED_LOOP1 MC_CLOSED_LOOP2	95
2.13.8	MC_THROTTLE_OFFSET0 MC_THROTTLE_OFFSET1 MC_THROTTLE_OFFSET2	95
2.13.9	MC_RANK_VIRTUAL_TEMPO MC_RANK_VIRTUAL_TEMP1 MC_RANK_VIRTUAL_TEMP2	96
2.13.10	MC_DDR_THERM_COMMAND0 MC_DDR_THERM_COMMAND1 MC_DDR_THERM_COMMAND2	96
2.13.11	MC_DDR_THERM_STATUS0 MC_DDR_THERM_STATUS1 MC_DDR_THERM_STATUS2	97
2.14	Integrated Memory Controller Miscellaneous Registers	97
2.14.1	MC_DIMM_CLK_RATIO_STATUS	97
2.14.2	MC_DIMM_CLK_RATIO	98



Tables


1-1	References.....	13
2-1	Functions Specifically Handled by the Processor	17
2-2	Device 0, Function 0: Generic Non-core Registers	18
2-3	Device 0, Function 1: System Address Decoder Registers	19
2-4	Device 2, Function 0: Intel QPI Link 0 Registers	20
2-5	Device 2, Function 1: Intel QPI Physical 0 Registers	21
2-6	Device 3, Function 0: Integrated Memory Controller Registers	22
2-7	Device 3, Function 1: Target Address Decoder Registers	23
2-8	Device 4, Function 0: Integrated Memory Controller Channel 0 Control Registers	24
2-9	Device 4, Function 1: Integrated Memory Controller Channel 0 Address Registers	25
2-10	Device 4, Function 2: Integrated Memory Controller Channel 0 Rank Registers	26
2-11	Device 4, Function 3: Integrated Memory Controller Channel 0 Thermal Control Registers.....	27
2-12	Device 5, Function 0: Integrated Memory Controller Channel 1 Control Registers	28
2-13	Device 5, Function 1: Integrated Memory Controller Channel 1 Address Registers	29
2-14	Device 5, Function 2: Integrated Memory Controller Channel 1 Rank Registers	30
2-15	Device 5, Function 3: Integrated Memory Controller Channel 1 Thermal Control Registers.....	31
2-16	Device 6, Function 0: Integrated Memory Controller Channel 2 Control Registers	32
2-17	Device 6, Function 1: Integrated Memory Controller Channel 2 Address Registers	33
2-18	Device 6, Function 2: Integrated Memory Controller Channel 2 Rank Registers	34
2-19	Device 6, Function 3: Integrated Memory Controller Channel 2 Thermal Control Registers.....	35



Revision History

Revision Number	Description	Date
001	Public release.	March 2009

§





1 Introduction

The Intel® Xeon® processor 3500 series are intended for Uni-processor (UP) workstation systems. The processor implements key new technologies:

- Integrated Memory Controller
- Point-to-point link interface based on Intel® QuickPath Interconnect (Intel® QPI). Reference to this interface may sometimes be abbreviated with Intel QPI throughout this document.

Note: In this document the Intel® Xeon® processor 3500 series will be referred to as "the processor."

This datasheet provides register descriptions for some of the registers located on the processor.

The processor is optimized for performance with the power efficiencies of a low-power microarchitecture to enable smaller, quieter systems.

The Intel Xeon processor 3500 Series are multi-core processors, based on 45 nm process technology. Processor features vary by component and include up to two Intel QuickPath Interconnect point to point links capable of up to 6.4 GT/s, up to 8 MB of shared cache, and an integrated memory controller. The processors support all the existing Streaming SIMD Extensions 2 (SSE2), Streaming SIMD Extensions 3 (SSE3) and Streaming SIMD Extensions 4 (SSE4). The processor supports several Advanced Technologies: Execute Disable Bit, Intel® 64 Technology, Enhanced Intel SpeedStep® Technology, Intel® Virtualization Technology (Intel® VT), Intel® Turbo Boost Technology, and Intel® Hyper-Threading Technology (Intel® HT Technology).

1.1 Terminology

A '#' symbol after a signal name refers to an active low signal, indicating a signal is in the active state when driven to a low level. For example, when RESET# is low, a reset has been requested.

1.1.1 Processor Terminology

Commonly used terms are explained here for clarification:

- **DDR3** — Double Data Rate 3 synchronous dynamic random access memory (SDRAM) is the name of the new DDR memory standard that is being developed as the successor to DDR2 SDRAM.
- **Enhanced Intel SpeedStep® Technology** — Enhanced Intel SpeedStep Technology allows trade-offs to be made between performance and power consumption.
- **Execute Disable Bit** — Execute Disable allows memory to be marked as executable or non-executable, when combined with a supporting operating system. If code attempts to run in non-executable memory the processor raises an error to the operating system. This feature can prevent some classes of viruses or worms that exploit buffer over run vulnerabilities and can thus help improve the overall security of the system. See the *Intel® 64 and IA-32 Architectures Software Developer's Manual* for more detailed information. Refer to <http://www.intel.com/> for future reference on up to date nomenclatures.



- **Eye Definitions** — The eye at any point along the data channel is defined to be the creation of overlapping of a large number of Unit Interval of the data signal and timing width measured with respect to the edges of a separate clock signal at any other point. Each differential signal pair by combining the D+ and D- signals produces a signal eye.
- **1366-Land LGA package** — The processor is available in a Flip-Chip Land Grid Array (FC-LGA) package, consisting of the processor die mounted on a land grid array substrate with an integrated heat spreader (IHS).
- **Functional Operation** — Refers to the normal operating conditions in which all processor specifications, including DC, AC, system bus, signal quality, mechanical, and thermal, are satisfied.
- **Integrated Memory Controller (IMC)** — A memory controller that is integrated in the processor silicon.
- **Integrated Heat Spreader (IHS)** — A component of the processor package used to enhance the thermal performance of the package. Component thermal solutions interface with the processor at the IHS surface.
- **Intel® 64 Architecture** — An enhancement to Intel's IA-32 architecture, allowing the processor to execute operating systems and applications written to take advantage of Intel 64. Further details on Intel 64 architecture and programming model can be found at <http://www.intel.com/technology/intel64/>.
- **Intel® QuickPath Interconnect** — A cache-coherent, link-based interconnect specification for Intel processor, chipset, and I/O bridge components. Sometimes abbreviated as Intel QPI.
- **Intel® QPI** — Abbreviation for Intel® QuickPath Interconnect.
- **Intel® Virtualization Technology (Intel® VT)** — A set of hardware enhancements to Intel server and client platforms that can improve virtualization solutions. Intel VT provides a foundation for widely-deployed virtualization solutions and enables more robust hardware assisted virtualization solutions. More information can be found at: <http://www.intel.com/technology/virtualization/>
- **Jitter** — Any timing variation of a transition edge or edges from the defined Unit Interval.
- **LGA1366 Socket** — The processor (in the LGA-1366 package) mates with the system board through this surface mount, 1366-contact socket.
- **Mirror Port** - Pads located on the top side of the processor package used to provide logic analyzer probing access for Intel QPI signal analysis.
- **Non-core** — The portion of the processor comprising the shared cache, IMC and Intel QPI Link interface.
- **OEM** — Original Equipment Manufacturer.
- **Storage Conditions** — Refers to 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 lands should not be connected to any supply voltages, have any I/Os biased, or receive any clocks.
- **Intel® Xeon® Processor 3500 Series** — The workstation product, including processor substrate and integrated heat spreader (IHS).



- **Unit Interval (UI)** — Signaling convention that is binary and unidirectional. In this binary signaling, one bit is sent for every edge of the forwarded clock, whether it be a rising edge or a falling edge. If a number of edges are collected at instances $t_1, t_2, t_3, \dots, t_k$ then the UI at instance "n" is defined as:

$$UI_n = t_n - t_{n-1}$$

1.2 References

Material and concepts available in the following documents may be beneficial when reading this document.


Table 1-1. References

Document	Document Number ¹
<i>Intel® Xeon® Processor 3500 Series Specification Update</i>	321335
<i>Intel® Xeon® Processor 3500 Series Datasheet, Volume 1</i>	321332
<i>Intel® 64 and IA-32 Intel® Architectures Software Developer's Manual</i>	
• <i>Volume 1: Basic Architecture</i>	253665
• <i>Volume 2A: Instruction Set Reference, A-M</i>	253666
• <i>Volume 2B: Instruction Set Reference, N-Z</i>	253667
• <i>Volume 3A: System Programming Guide, Part 1</i>	253668
• <i>Volume 3B: Systems Programming Guide, Part 2</i>	253669

Note:

1. Documents are available publicly at <http://www.intel.com>.

§



2 Register Description

The processor supports PCI configuration space accesses using the mechanism denoted as Configuration Mechanism in the PCI specification as defined in the *PCI Local Bus Specification*, Revision 2.3, as well as the PCI Express* enhanced configuration mechanism as specified in the *PCI Express Base Specification*, Revision 1.1. All the registers are organized by bus, device, function, etc. as defined in the *PCI Express Base Specification*, Revision 1.1. All processor registers appear on the PCI bus assigned for the processor socket. Bus number is derived by the max bus range setting and processor socket number. All multi-byte numeric fields use “little-endian” ordering (i.e., lower addresses contain the least significant parts of the field).

As processor features vary by component, not all of the register descriptions in this document apply to all processors. This document highlights registers which do not apply to all processor components. Refer to the particular processor's Specification Update for a list of features supported.

2.1 Register Terminology

Registers and register bits are assigned one or more of the following attributes. These attributes define the behavior of register and the bit(s) that are contained within. All bits are set to default values by hard reset. Sticky bits retain their states between hard resets.

Term	Description
RO	Read Only. If a register bit is read only, the hardware sets its state. The bit may be read by software. Writes to this bit have no effect.
WO	Write Only. The register bit is not implemented as a bit. The write causes some hardware event to take place.
RW	Read/Write. A register bit with this attribute can be read and written by software.
RC	Read Clear: The bit or bits can be read by software, but the act of reading causes the value to be cleared.
RCW	Read Clear/Write: A register bit with this attribute will get cleared after the read. The register bit can be written.
RW1C	Read/Write 1 Clear. A register bit with this attribute can be read or cleared by software. In order to clear this bit, a one must be written to it. Writing a zero will have no effect.
RWOC	Read/Write 0 Clear. A register bit with this attribute can be read or cleared by software. In order to clear this bit, a zero must be written to it. Writing a one will have no effect.
RW1S	Read/Write 1 Set: A register bit can be either read or set by software. In order to set this bit, a one must be written to it. Writing a zero to this bit has no effect. Hardware will clear this bit.
RWOS	Read/Write 0 Set: A register bit can be either read or set by software. In order to set this bit, a zero must be written to it. Writing a one to this bit has no effect. Hardware will clear this bit.
RWL	Read/Write/Lock. A register bit with this attribute can be read or written by software. Hardware or a configuration bit can lock the bit and prevent it from being updated.
RWO	Read/Write Once. A register bit with this attribute can be written to only once after power up. After the first write, the bit becomes read only. This attribute is applied on a bit by bit basis. For example, if the RWO attribute is applied to a 2 bit field, and only one bit is written, then the written bit cannot be rewritten (unless reset). The unwritten bit, of the field, may still be written once. This is special case of RWL.
RRW	Read/Restricted Write. This bit can be read and written by software. However, only supported values will be written. Writes of non supported values will have no effect.
L	Lock. A register bit with this attribute becomes Read Only after a lock bit is set.

Term	Description
RSVD	Reserved Bit. This bit is reserved for future expansion and must not be written. The <i>PCI Local Bus Specification</i> , Revision 2.2 requires that reserved bits must be preserved. Any software that modifies a register that contains a reserved bit is responsible for reading the register, modifying the desired bits, and writing back the result.
Reserved Bits	Some of the processor registers described in this section contain reserved bits. These bits are labeled "Reserved". Software must deal correctly with fields that are reserved. On reads, software must use appropriate masks to extract the defined bits and not rely on reserved bits being any particular value. On writes, software must ensure that the values of reserved bit positions are preserved. That is, the values of reserved bit positions must first be read, merged with the new values for other bit positions and then written back. Note that software does not need to perform a read-merge-write operation for the Configuration Address (CONFIG_ADDRESS) register.
Reserved Registers	In addition to reserved bits within a register, the processor contains address locations in the configuration space that are marked either "Reserved" or "Intel Reserved". The processor responds to accesses to "Reserved" address locations by completing the host cycle. When a "Reserved" register location is read, a zero value is returned. ("Reserved" registers can be 8, 16, or 32 bits in size). Writes to "Reserved" registers have no effect on the processor. Registers that are marked as "Intel Reserved" must not be modified by system software. Writes to "Intel Reserved" registers may cause system failure. Reads to "Intel Reserved" registers may return a non-zero value.
Default Value upon a Reset	Upon a reset, the processor sets all of its internal configuration registers to predetermined default states. Some register values at reset are determined by external strapping options. The default state represents the minimum functionality feature set required to successfully bring up the system. Hence, it does not represent the optimal system configuration. It is the responsibility of the system initialization software (usually BIOS) to properly determine the DRAM configurations, operating parameters and optional system features that are applicable, and to program the processor registers accordingly.
"ST" appended to the end of a bit name	The bit is "sticky" or unchanged by a hard reset. These bits can only be cleared by a PWRGOOD reset.

2.2 Platform Configuration Structure

The processor contains 6 PCI devices within a single physical component. The configuration registers for these devices are mapped as devices residing on the PCI bus assigned for the processor socket. Bus number is derived by the max bus range setting and processor socket number.

- **Device 0:** Generic processor non-core. Device 0, Function 0 contains the generic non-core configuration registers for the processor and resides at DID (Device ID) of 2C41h. Device 0, Function 1 contains the System Address Decode registers and resides at DID of 2C01h.
- **Device 2:** Intel QPI. Device 2, Function 0 contains the Intel® QuickPath Interconnect configuration registers for Intel QPI Link 0 and resides at DID of 2C10h. Device 2, Function 1 contains the physical layer registers for Intel QPI Link 0 and resides at DID of 2C11h.
- **Device 3:** Integrated Memory Controller. Device 3, Function 0 contains the general registers for the Integrated Memory Controller and resides at DID of 2C18h. Device 3, Function 1 contains the Target Address Decode registers for the Integrated Memory Controller and resides at DID of 2C19h. Device 3, Function 2 contains the RAS registers for the Integrated Memory Controller and resides at DID of 2C1Ah. Device 3, Function 4 contains the test registers for the Integrated Memory Controller and resides at DID of 2C1Ch. Function 2 only applies to processors supporting registered DIMMs.
- **Device 4:** Integrated Memory Controller Channel 0. Device 4, Function 0 contains the control registers for Integrated Memory Controller Channel 0 and resides at DID of 2C20h. Device 4, Function 1 contains the address registers for Integrated Memory Controller Channel 0 and resides at DID of 2C21h. Device 4, Function 2 contains the rank registers for Integrated Memory Controller Channel 0 and resides

at DID of 2C22h. Device 4, Function 3 contains the thermal control registers for Integrated Memory Controller Channel 0 and resides at DID of 2C23h.

- **Device 5:** Integrated Memory Controller Channel 1. Device 5, Function 0 contains the control registers for Integrated Memory Controller Channel 1 and resides at DID of 2C28h. Device 5, Function 1 contains the address registers for Integrated Memory Controller Channel 1 and resides at DID of 2C29h. Device 5, Function 2 contains the rank registers for Integrated Memory Controller Channel 1 and resides at DID of 2C2Ah. Device 5, Function 3 contains the thermal control registers for Integrated Memory Controller Channel 1 and resides at DID of 2C2Bh.
- **Device 6:** Integrated Memory Controller Channel 2. Device 6, Function 0 contains the control registers for Integrated Memory Controller Channel 2 and resides at DID of 2C30h. Device 6, Function 1 contains the address registers for Integrated Memory Controller Channel 2 and resides at DID of 2C31h. Device 6, Function 2 contains the rank registers for Integrated Memory Controller Channel 2 and resides at DID of 2C32h. Device 6, Function 3 contains the thermal control registers for Integrated Memory Controller Channel 2 and resides at DID of 2C33h.

2.3 Device Mapping

Each component in the processor is uniquely identified by a PCI bus address consisting of Bus Number, Device Number, and Function Number. Device configuration is based on the PCI Type 0 configuration conventions. All processor registers appear on the PCI bus assigned for the processor socket. Bus number is derived by the max bus range setting and processor socket number.

Table 2-1. Functions Specifically Handled by the Processor

Component	Register Group	DID	Device	Function
Processor	Intel QuickPath Architecture Generic Non-core Registers	2C41h	0	0
	Intel QuickPath Architecture System Address Decoder	2C01h		1
	Intel QPI Link 0	2C10h	2	0
	Intel QPI Physical 0	2C11		1
	Integrated Memory Controller Registers	2C18h	3	0
	Integrated Memory Controller Target Address Decoder	2C19h		1
	Integrated Memory Controller RAS Registers	2C1Ah		2 ¹
	Integrated Memory Controller Test Registers	2C1Ch		4
	Integrated Memory Controller Channel 0 Control	2C20h	4	0
	Integrated Memory Controller Channel 0 Address	2C21h		1
	Integrated Memory Controller Channel 0 Rank	2C22h		2
	Integrated Memory Controller Channel 0 Thermal Control	2C23h		3
	Integrated Memory Controller Channel 1 Control	2C28h	5	0
	Integrated Memory Controller Channel 1 Address	2C29h		1
	Integrated Memory Controller Channel 1 Rank	2C2Ah		2
	Integrated Memory Controller Channel 1 Thermal Control	2C2Bh		3
	Integrated Memory Controller Channel 2 Control	2C30h	6	0
	Integrated Memory Controller Channel 2 Address	2C31h		1
	Integrated Memory Controller Channel 2 Rank	2C32h		2
	Integrated Memory Controller Channel 2 Thermal Control	2C33h		3

Notes:

2. Applies only to processors supporting sparing, mirroring, and scrubbing RAS features.

2.4 Detailed Configuration Space Maps

Table 2-2. Device 0, Function 0: Generic Non-core Registers

DID	VID		
PCISTS	PCICMD		
CCR	RID		
HDR			
SID	SVID		
		00h	80h
		04h	84h
		08h	88h
		0Ch	8Ch
		10h	90h
		14h	94h
		18h	98h
		1Ch	9Ch
		20h	A0h
		24h	A4h
		28h	A8h
		2Ch	ACh
		30h	B0h
		34h	B4h
		38h	B8h
		3Ch	BCh
		40h	C0h
		44h	C4h
		48h	C8h
		4Ch	CCh
		50h	D0h
		54h	D4h
		58h	D8h
		5Ch	DCh
		60h	E0h
		64h	E4h
		68h	E8h
		6Ch	EC
		70h	F0h
		74h	F4h
		78h	F8h
		7Ch	FC

Table 2-3. Device 0, Function 1: System Address Decoder Registers

DID	VID		
PCISTS	PCICMD		
CCR	RID		
	HDR		
SID	SVID		
SAD_PAM0123			
SAD_PAM456			
SAD_HEN			
SAD_SMRAM			
SAD_PCIEBAR			
00h	SAD_DRAM_RULE_0		80h
04h	SAD_DRAM_RULE_1		84h
08h	SAD_DRAM_RULE_2		88h
0Ch	SAD_DRAM_RULE_3		8Ch
10h	SAD_DRAM_RULE_4		90h
14h	SAD_DRAM_RULE_5		94h
18h	SAD_DRAM_RULE_6		98h
1Ch	SAD_DRAM_RULE_7		9Ch
20h			A0h
24h			A4h
28h			A8h
2Ch			ACh
30h			B0h
34h			B4h
38h			B8h
3Ch			BCh
40h	SAD_INTERLEAVE_LIST_0		C0h
44h	SAD_INTERLEAVE_LIST_1		C4h
48h	SAD_INTERLEAVE_LIST_2		C8h
4Ch	SAD_INTERLEAVE_LIST_3		CCh
50h	SAD_INTERLEAVE_LIST_4		D0h
54h	SAD_INTERLEAVE_LIST_5		D4h
58h	SAD_INTERLEAVE_LIST_6		D8h
5Ch	SAD_INTERLEAVE_LIST_7		DCh
60h			E0h
64h			E4h
68h			E8h
6Ch			ECh
70h			F0h
74h			F4h
78h			F8h
7Ch			FCh

Table 2-4. Device 2, Function 0: Intel QPI Link 0 Registers

DID	VID		
PCISTS	PCICMD		
CCR	RID		
	HDR		
SID	SVID		
QPI_QPILCL_L0			
		00h	80h
		04h	84h
		08h	88h
		0Ch	8Ch
		10h	90h
		14h	94h
		18h	98h
		1Ch	9Ch
		20h	A0h
		24h	A4h
		28h	A8h
		2Ch	ACh
		30h	B0h
		34h	B4h
		38h	B8h
		3Ch	BCh
		40h	C0h
		44h	C4h
		48h	C8h
		4Ch	CCh
		50h	D0h
		54h	D4h
		58h	D8h
		5Ch	DCh
		60h	E0h
		64h	E4h
		68h	E8h
		6Ch	ECh
		70h	F0h
		74h	F4h
		78h	F8h
		7Ch	FCCh

Table 2-5. Device 2, Function 1: Intel QPI Physical 0 Registers

DID	VID		80h
PCISTS	PCICMD		84h
CCR	RID		88h
HDR			8Ch
			90h
			94h
			98h
			9Ch
			A0h
			A4h
			A8h
SID	SVID		ACh
			B0h
			B4h
			B8h
			BCh
			C0h
			C4h
			C8h
			CCh
			D0h
			D4h
			D8h
			DCh
			E0h
			E4h
			E8h
			EC
			F0h
			F4h
			F8h
			FC



Table 2-6. Device 3, Function 0: Integrated Memory Controller Registers

DID	VID		80h
PCISTS	PCICMD		84h
CCR	RID		88h
HDR			8Ch
			90h
			94h
			98h
			9Ch
			A0h
			A4h
			A8h
SID	SVID		ACh
			B0h
			B4h
			B8h
			BCh
			C0h
			C4h
MC_CONTROL			C8h
MC_STATUS			CCh
MC_SMI_SPARE_DIMM_ERROR_STATUS			D0h
MC_SMI_SPARE_CNTRL			D4h
			D8h
MC_RESET_CONTROL			DCh
MC_CHANNEL_MAPPER			E0h
MC_MAX_DOD			E4h
			E8h
			ECCh
MC_RD_CRDT_INIT			F0h
MC_CRDT_WR_THLD			F4h
MC_SCRUBADDR_LO			F8h
MC_SCRUBADDR_HI			FCh

Table 2-7. Device 3, Function 1: Target Address Decoder Registers

Table 2-8. Device 4, Function 0: Integrated Memory Controller Channel 0 Control Registers

DID	VID		
PCISTS	PCICMD		
CCR	RID		
	HDR		
SID	SVID		
MC_CHANNEL_0_DIMM_RESET_CMD			
MC_CHANNEL_0_DIMM_INIT_CMD			
MC_CHANNEL_0_DIMM_INIT_PARAMS			
MC_CHANNEL_0_DIMM_INIT_STATUS			
MC_CHANNEL_0_DDR3CMD			
MC_CHANNEL_0_REFRESH_THROTTLE_SUPPORT			
MC_CHANNEL_0_MRS_VALUE_0_1			
MC_CHANNEL_0_MRS_VALUE_2			
MC_CHANNEL_0_RANK_PRESENT			
00h	MC_CHANNEL_0_RANK_TIMING_A	80h	
04h	MC_CHANNEL_0_RANK_TIMING_B	84h	
08h	MC_CHANNEL_0_BANK_TIMING	88h	
0Ch	MC_CHANNEL_0_REFRESH_TIMING	8Ch	
10h	MC_CHANNEL_0_CKE_TIMING	90h	
14h	MC_CHANNEL_0_ZQ_TIMING	94h	
18h	MC_CHANNEL_0_RCOMP_PARAMS	98h	
1Ch	MC_CHANNEL_0_ODT_PARAMS1	9Ch	
20h	MC_CHANNEL_0_ODT_PARAMS2	A0h	
24h	MC_CHANNEL_0_ODT_MATRIX_RANK_0_3_RD	A4h	
28h	MC_CHANNEL_0_ODT_MATRIX_RANK_4_7_RD	A8h	
2Ch	MC_CHANNEL_0_ODT_MATRIX_RANK_0_3_WR	ACh	
30h	MC_CHANNEL_0_ODT_MATRIX_RANK_4_7_WR	B0h	
34h	MC_CHANNEL_0_WAO_PARAMS	B4h	
38h	MC_CHANNEL_0_SCHEDULER_PARAMS	B8h	
3Ch	MC_CHANNEL_0_MAINTENANCE_OPS	BCh	
40h	MC_CHANNEL_0_TX_BG_SETTINGS	C0h	
44h		C4h	
48h	MC_CHANNEL_0_RX_BGF_SETTINGS	C8h	
4Ch	MC_CHANNEL_0_EW_BGF_SETTINGS	CCh	
50h	MC_CHANNEL_0_EW_BGF_OFFSET_SETTINGS	D0h	
54h	MC_CHANNEL_0_ROUND_TRIP_LATENCY	D4h	
58h	MC_CHANNEL_0_PAGETABLE_PARAMS1	D8h	
5Ch	MC_CHANNEL_0_PAGETABLE_PARAMS2	DCh	
60h	MC_TX_BG_CMD_DATA_RATIO_SETTING_CHO	E0h	
64h	MC_TX_BG_CMD_OFFSET_SETTINGS_CHO	E4h	
68h	MC_TX_BG_DATA_OFFSET_SETTINGS_CHO	E8h	
6Ch		ECh	
70h	MC_CHANNEL_0_ADDR_MATCH	F0h	
74h		F4h	
78h	MC_CHANNEL_0_ECC_ERROR_MASK	F8h	
7Ch	MC_CHANNEL_0_ECC_ERROR_INJECT	FCh	

Table 2-9. Device 4, Function 1: Integrated Memory Controller Channel 0 Address Registers

Table 2-10. Device 4, Function 2: Integrated Memory Controller Channel 0 Rank Registers

DID	VID		
PCISTS	PCICMD		
CCR	RID		
	HDR		
SID	SVID		
MC_RIR_LIMIT_CHO_0			
MC_RIR_LIMIT_CHO_1			
MC_RIR_LIMIT_CHO_2			
MC_RIR_LIMIT_CHO_3			
MC_RIR_LIMIT_CHO_4			
MC_RIR_LIMIT_CHO_5			
MC_RIR_LIMIT_CHO_6			
MC_RIR_LIMIT_CHO_7			
00h	MC_RIR_WAY_CHO_0	80h	
04h	MC_RIR_WAY_CHO_1	84h	
08h	MC_RIR_WAY_CHO_2	88h	
0Ch	MC_RIR_WAY_CHO_3	8Ch	
10h	MC_RIR_WAY_CHO_4	90h	
14h	MC_RIR_WAY_CHO_5	94h	
18h	MC_RIR_WAY_CHO_6	98h	
1Ch	MC_RIR_WAY_CHO_7	9Ch	
20h	MC_RIR_WAY_CHO_8	A0h	
24h	MC_RIR_WAY_CHO_9	A4h	
28h	MC_RIR_WAY_CHO_10	A8h	
2Ch	MC_RIR_WAY_CHO_11	ACh	
30h	MC_RIR_WAY_CHO_12	B0h	
34h	MC_RIR_WAY_CHO_13	B4h	
38h	MC_RIR_WAY_CHO_14	B8h	
3Ch	MC_RIR_WAY_CHO_15	BCh	
40h	MC_RIR_WAY_CHO_16	C0h	
44h	MC_RIR_WAY_CHO_17	C4h	
48h	MC_RIR_WAY_CHO_18	C8h	
4Ch	MC_RIR_WAY_CHO_19	CCh	
50h	MC_RIR_WAY_CHO_20	D0h	
54h	MC_RIR_WAY_CHO_21	D4h	
58h	MC_RIR_WAY_CHO_22	D8h	
5Ch	MC_RIR_WAY_CHO_23	DCh	
60h	MC_RIR_WAY_CHO_24	E0h	
64h	MC_RIR_WAY_CHO_25	E4h	
68h	MC_RIR_WAY_CHO_26	E8h	
6Ch	MC_RIR_WAY_CHO_27	ECh	
70h	MC_RIR_WAY_CHO_28	F0h	
74h	MC_RIR_WAY_CHO_29	F4h	
78h	MC_RIR_WAY_CHO_30	F8h	
7Ch	MC_RIR_WAY_CHO_31	FCh	

Table 2-11. Device 4, Function 3: Integrated Memory Controller Channel 0 Thermal Control Registers

Table 2-12. Device 5, Function 0: Integrated Memory Controller Channel 1 Control Registers

DID	VID		
PCISTS	PCICMD		
CCR	RID		
	HDR		
SID	SVID		
MC_CHANNEL_1_DIMM_RESET_CMD			
MC_CHANNEL_1_DIMM_INIT_CMD			
MC_CHANNEL_1_DIMM_INIT_PARAMS			
MC_CHANNEL_1_DIMM_INIT_STATUS			
MC_CHANNEL_1_DDR3CMD			
MC_CHANNEL_1_REFRESH_THROTTLE_SUPPORT			
MC_CHANNEL_1_MRS_VALUE_0_1			
MC_CHANNEL_1_MRS_VALUE_2			
MC_CHANNEL_1_RANK_PRESENT			
00h	MC_CHANNEL_1_RANK_TIMING_A	80h	
04h	MC_CHANNEL_1_RANK_TIMING_B	84h	
08h	MC_CHANNEL_1_BANK_TIMING	88h	
0Ch	MC_CHANNEL_1_REFRESH_TIMING	8Ch	
10h	MC_CHANNEL_1_CKE_TIMING	90h	
14h	MC_CHANNEL_1_ZQ_TIMING	94h	
18h	MC_CHANNEL_1_RCOMP_PARAMS	98h	
1Ch	MC_CHANNEL_1_ODT_PARAMS1	9Ch	
20h	MC_CHANNEL_1_ODT_PARAMS2	A0h	
24h	MC_CHANNEL_1_ODT_MATRIX_RANK_0_3_RD	A4h	
28h	MC_CHANNEL_1_ODT_MATRIX_RANK_4_7_RD	A8h	
2Ch	MC_CHANNEL_1_ODT_MATRIX_RANK_0_3_WR	ACh	
30h	MC_CHANNEL_1_ODT_MATRIX_RANK_4_7_WR	B0h	
34h	MC_CHANNEL_1_WAO_PARAMS	B4h	
38h	MC_CHANNEL_1_SCHEDULER_PARAMS	B8h	
3Ch	MC_CHANNEL_1_MAINTENANCE_OPS	BCh	
40h	MC_CHANNEL_1_TX_BG_SETTINGS	C0h	
44h		C4h	
48h	MC_CHANNEL_1_RX_BGF_SETTINGS	C8h	
4Ch	MC_CHANNEL_1_EW_BGF_SETTINGS	CCh	
50h	MC_CHANNEL_1_EW_BGF_OFFSET_SETTINGS	D0h	
54h	MC_CHANNEL_1_ROUND_TRIP_LATENCY	D4h	
58h	MC_CHANNEL_1_PAGETABLE_PARAMS1	D8h	
5Ch	MC_CHANNEL_1_PAGETABLE_PARAMS2	DCh	
60h	MC_TX_BG_CMD_DATA_RATIO_SETTING_CH1	E0h	
64h	MC_TX_BG_CMD_OFFSET_SETTINGS_CH1	E4h	
68h	MC_TX_BG_DATA_OFFSET_SETTINGS_CH1	E8h	
6Ch		ECh	
70h	MC_CHANNEL_1_ADDR_MATCH	F0h	
74h		F4h	
78h	MC_CHANNEL_1_ECC_ERROR_MASK	F8h	
7Ch	MC_CHANNEL_1_ECC_ERROR_INJECT	FCh	

Table 2-13. Device 5, Function 1: Integrated Memory Controller Channel 1 Address Registers

Table 2-14. Device 5, Function 2: Integrated Memory Controller Channel 1 Rank Registers

DID	VID	
PCISTS	PCICMD	
CCR	RID	
	HDR	
SID	SVID	
MC_RIR_LIMIT_CH1_0		
MC_RIR_LIMIT_CH1_1		
MC_RIR_LIMIT_CH1_2		
MC_RIR_LIMIT_CH1_3		
MC_RIR_LIMIT_CH1_4		
MC_RIR_LIMIT_CH1_5		
MC_RIR_LIMIT_CH1_6		
MC_RIR_LIMIT_CH1_7		
MC_RIR_LIMIT_CH1_8		
MC_RIR_LIMIT_CH1_9		
MC_RIR_LIMIT_CH1_10		
MC_RIR_LIMIT_CH1_11		
MC_RIR_LIMIT_CH1_12		
MC_RIR_LIMIT_CH1_13		
MC_RIR_LIMIT_CH1_14		
MC_RIR_LIMIT_CH1_15		
MC_RIR_LIMIT_CH1_16		
MC_RIR_LIMIT_CH1_17		
MC_RIR_LIMIT_CH1_18		
MC_RIR_LIMIT_CH1_19		
MC_RIR_LIMIT_CH1_20		
MC_RIR_LIMIT_CH1_21		
MC_RIR_LIMIT_CH1_22		
MC_RIR_LIMIT_CH1_23		
MC_RIR_LIMIT_CH1_24		
MC_RIR_LIMIT_CH1_25		
MC_RIR_LIMIT_CH1_26		
MC_RIR_LIMIT_CH1_27		
MC_RIR_LIMIT_CH1_28		
MC_RIR_LIMIT_CH1_29		
MC_RIR_LIMIT_CH1_30		
MC_RIR_LIMIT_CH1_31		

Table 2-15. Device 5, Function 3: Integrated Memory Controller Channel 1 Thermal Control Registers

Table 2-16. Device 6, Function 0: Integrated Memory Controller Channel 2 Control Registers

DID	VID		
PCISTS	PCICMD		
CCR	RID		
	HDR		
SID	SVID		
MC_CHANNEL_2_DIMM_RESET_CMD			
MC_CHANNEL_2_DIMM_INIT_CMD			
MC_CHANNEL_2_DIMM_INIT_PARAMS			
MC_CHANNEL_2_DIMM_INIT_STATUS			
MC_CHANNEL_2_DDR3CMD			
MC_CHANNEL_2_REFRESH_THROTTLE_SUPPORT			
MC_CHANNEL_2_MRS_VALUE_0_1			
MC_CHANNEL_2_MRS_VALUE_2			
MC_CHANNEL_2_RANK_PRESENT			
00h	MC_CHANNEL_2_RANK_TIMING_A	80h	
04h	MC_CHANNEL_2_RANK_TIMING_B	84h	
08h	MC_CHANNEL_2_BANK_TIMING	88h	
0Ch	MC_CHANNEL_2_REFRESH_TIMING	8Ch	
10h	MC_CHANNEL_2_CKE_TIMING	90h	
14h	MC_CHANNEL_2_ZQ_TIMING	94h	
18h	MC_CHANNEL_2_RCOMP_PARAMS	98h	
1Ch	MC_CHANNEL_2_ODT_PARAMS1	9Ch	
20h	MC_CHANNEL_2_ODT_PARAMS2	A0h	
24h	MC_CHANNEL_2_ODT_MATRIX_RANK_0_3_RD	A4h	
28h	MC_CHANNEL_2_ODT_MATRIX_RANK_4_7_RD	A8h	
2Ch	MC_CHANNEL_2_ODT_MATRIX_RANK_0_3_WR	ACh	
30h	MC_CHANNEL_2_ODT_MATRIX_RANK_4_7_WR	B0h	
34h	MC_CHANNEL_2_WAO_PARAMS	B4h	
38h	MC_CHANNEL_2_SCHEDULER_PARAMS	B8h	
3Ch	MC_CHANNEL_2_MAINTENANCE_OPS	BCh	
40h	MC_CHANNEL_2_TX_BG_SETTINGS	C0h	
44h		C4h	
48h	MC_CHANNEL_2_RX_BGF_SETTINGS	C8h	
4Ch	MC_CHANNEL_2_EW_BGF_SETTINGS	CCh	
50h	MC_CHANNEL_2_EW_BGF_OFFSET_SETTINGS	D0h	
54h	MC_CHANNEL_2_ROUND_TRIP_LATENCY	D4h	
58h	MC_CHANNEL_2_PAGETABLE_PARAMS1	D8h	
5Ch	MC_CHANNEL_2_PAGETABLE_PARAMS2	DCh	
60h	MC_TX_BG_CMD_DATA_RATIO_SETTING_CH2	E0h	
64h	MC_TX_BG_CMD_OFFSET_SETTINGS_CH2	E4h	
68h	MC_TX_BG_DATA_OFFSET_SETTINGS_CH2	E8h	
6Ch		ECh	
70h	MC_CHANNEL_2_ADDR_MATCH	F0h	
74h		F4h	
78h	MC_CHANNEL_2_ECC_ERROR_MASK	F8h	
7Ch	MC_CHANNEL_2_ECC_ERROR_INJECT	FCh	

Table 2-17. Device 6, Function 1: Integrated Memory Controller Channel 2 Address Registers

Table 2-18. Device 6, Function 2: Integrated Memory Controller Channel 2 Rank Registers

DID	VID		
PCISTS	PCICMD		
CCR	RID		
	HDR		
SID	SVID		
MC_RIR_LIMIT_CH2_0			
MC_RIR_LIMIT_CH2_1			
MC_RIR_LIMIT_CH2_2			
MC_RIR_LIMIT_CH2_3			
MC_RIR_LIMIT_CH2_4			
MC_RIR_LIMIT_CH2_5			
MC_RIR_LIMIT_CH2_6			
MC_RIR_LIMIT_CH2_7			
00h	MC_RIR_WAY_CH2_0	80h	
04h	MC_RIR_WAY_CH2_1	84h	
08h	MC_RIR_WAY_CH2_2	88h	
0Ch	MC_RIR_WAY_CH2_3	8Ch	
10h	MC_RIR_WAY_CH2_4	90h	
14h	MC_RIR_WAY_CH2_5	94h	
18h	MC_RIR_WAY_CH2_6	98h	
1Ch	MC_RIR_WAY_CH2_7	9Ch	
20h	MC_RIR_WAY_CH2_8	A0h	
24h	MC_RIR_WAY_CH2_9	A4h	
28h	MC_RIR_WAY_CH2_10	A8h	
2Ch	MC_RIR_WAY_CH2_11	ACh	
30h	MC_RIR_WAY_CH2_12	B0h	
34h	MC_RIR_WAY_CH2_13	B4h	
38h	MC_RIR_WAY_CH2_14	B8h	
3Ch	MC_RIR_WAY_CH2_15	BCh	
40h	MC_RIR_WAY_CH2_16	C0h	
44h	MC_RIR_WAY_CH2_17	C4h	
48h	MC_RIR_WAY_CH2_18	C8h	
4Ch	MC_RIR_WAY_CH2_19	CCh	
50h	MC_RIR_WAY_CH2_20	D0h	
54h	MC_RIR_WAY_CH2_21	D4h	
58h	MC_RIR_WAY_CH2_22	D8h	
5Ch	MC_RIR_WAY_CH2_23	DCh	
60h	MC_RIR_WAY_CH2_24	E0h	
64h	MC_RIR_WAY_CH2_25	E4h	
68h	MC_RIR_WAY_CH2_26	E8h	
6Ch	MC_RIR_WAY_CH2_27	ECh	
70h	MC_RIR_WAY_CH2_28	F0h	
74h	MC_RIR_WAY_CH2_29	F4h	
78h	MC_RIR_WAY_CH2_30	F8h	
7Ch	MC_RIR_WAY_CH2_31	FCh	

Table 2-19. Device 6, Function 3: Integrated Memory Controller Channel 2 Thermal Control Registers

2.5 PCI Standard Registers

These registers appear in every function for every device.

Note: Reserved bit locations are not shown in the following register tables.

2.5.1 VID - Vendor Identification Register

The VID Register contains the vendor identification number. This 16-bit register, combined with the Device Identification Register uniquely identifies the manufacturer of the function within the processor. Writes to this register have no effect.

Device:	0		
Function:	0-1		
Offset:	00h		
Device:	2		
Function:	0-1, 4-5		
Offset:	00h		
Device:	3		
Function:	0-2, 4		
Offset:	00h		
Device:	4-6		
Function:	0-3		
Offset:	00h		
Bit	Type	Reset Value	Description
15:0	RO	8086h	Vendor Identification Number The value assigned to Intel.

2.5.2 DID - Device Identification Register

This 16-bit register combined with the Vendor Identification register uniquely identifies the Function within the processor. Writes to this register have no effect. See [Table 2-1](#) for the DID of each processor function.

Device:	0		
Function:	0-1		
Offset:	02h		
Device:	2		
Function:	0-1, 4-5		
Offset:	02h		
Device:	3		
Function:	0-2, 4		
Offset:	02h		
Device:	4-6		
Function:	0-3		
Offset:	02h		
Bit	Type	Reset Value	Description
15:0	RO	*See Table 2-1	Device Identification Number Identifies each function of the processor.

2.5.3 RID - Revision Identification Register

This register contains the revision number of the processor. The Revision ID (RID) is a traditional 8-bit Read Only (RO) register located at offset 08h in the standard PCI header of every PCI/PCI Express compatible device and function.

Device: 0	Function: 0-1	Offset: 08h	
Device: 2	Function: 0-1, 4-5	Offset: 08h	
Device: 3	Function: 0-2, 4	Offset: 08h	
Device: 4-6	Function: 0-3	Offset: 08h	
Bit	Type	Reset Value	Description
7:0	RO	0h	Revision Identification Number Refer to the <i>Intel® Xeon® Processor 3500 Series Specification Update</i> for the value of the Revision ID Register.

2.5.4 CCR - Class Code Register

This register contains the Class Code for the device. Writes to this register have no effect.

Device: 0	Function: 0-1	Offset: 09h	
Device: 2	Function: 0-1, 4-5	Offset: 09h	
Device: 3	Function: 0-2, 4	Offset: 09h	
Device: 4-6	Function: 0-3	Offset: 09h	
Bit	Type	Reset Value	Description
23:16	RO	06h	Base Class This field indicates the general device category. For the processor, this field is hardwired to 06h, indicating it is a "Bridge Device".
15:8	RO	0	Sub-Class This field qualifies the Base Class, providing a more detailed specification of the device function. For all devices the default is 00h, indicating "Host Bridge".
7:0	RO	0	Register-Level Programming Interface This field identifies a specific programming interface (if any), that device independent software can use to interact with the device. There are no such interfaces defined for "Host Bridge" types, and this field is hardwired to 00h.

2.5.5 HDR - Header Type Register

This register identifies the header layout of the configuration space.

Device: 0	Function: 0-1	Offset: 0Eh	
Device: 2	Function: 0-1, 4-5	Offset: 0Eh	
Device: 3	Function: 0-2, 4	Offset: 0Eh	
Device: 4-6	Function: 0-3	Offset: 0Eh	
Bit	Type	Reset Value	Description
7	RO	1	Multi-function Device Selects whether this is a multi-function device, that may have alternative configuration layouts. This bit is hardwired to 1 for devices in the processor.
6:0	RO	0	Configuration Layout This field identifies the format of the configuration header layout for a PCI-to-PCI bridge from bytes 10h through 3Fh. For all devices the default is 00h, indicating a conventional type 00h PCI header.

2.5.6 SID/SVID - Subsystem Identity/Subsystem Vendor Identification Register

This register identifies the manufacturer of the system. This 32-bit register uniquely identifies any PCI device.

Device: 0	Function: 0-1	Offset: 2Ch, 2Eh	
Device: 2	Function: 0-1, 4-5	Offset: 2Ch, 2Eh	
Device: 3	Function: 0-2, 4	Offset: 2Ch, 2Eh	
Device: 4-6	Function: 0-3	Offset: 2Ch, 2Eh	
Access as a Dword			
Bit	Type	Reset Value	Description
31:16	RWO	8086h	Subsystem Identification Number The default value specifies Intel
15:0	RWO	8086h	Vendor Identification Number The default value specifies Intel.

2.5.7 PCI CMD - Command Register

This register defines the PCI 3.0 compatible command register values applicable to PCI Express space.

Device: 0	Function: 0-1	Offset: 04h	
Device: 2	Function: 0-1, 4-5	Offset: 04h	
Device: 3	Function: 0-2, 4	Offset: 04h	
Device: 4-6	Function: 0-3	Offset: 04h	
Bit Type Reset Value Description			
15:11	RV	0	Reserved. (by PCI SIG)
10	RO	0	<p>INTxDisable: Interrupt Disable Controls the ability of the PCI Express port to generate INTx messages. If this device does not generate interrupts then this bit is not implemented and is RO.</p> <p>If this device generates interrupts then this bit is RW and this bit disables the device/function from asserting INTx#. A value of 0 enables the assertion of its INTx# signal. A value of 1 disables the assertion of its INTx# signal.</p> <p>1 = Legacy Interrupt mode is disabled 0 = Legacy Interrupt mode is enabled</p>
9	RO	0	<p>FB2B: Fast Back-to-Back Enable This bit controls whether or not the master can do fast back-to-back writes. Since this device is strictly a target this bit is not implemented. This bit is hardwired to 0. Writes to this bit position have no effect.</p>
8	RO	0	<p>SERRE: SERR Message Enable This bit is a global enable bit for this device's SERR messaging. This host bridge will not implement SERR messaging. This bit is hardwired to 0. If SERR is used for error generation, then this bit must be RW and enable/disable SERR signaling.</p>
7	RO	0	<p>IDSELWCC: IDSEL Stepping/Wait Cycle Control Per PCI 2.3 specification this bit is hardwired to 0.</p>
6	RO	0	<p>PERRE: Parity Error Response Enable Parity error is not implemented in this host bridge. This bit is hardwired to 0.</p>
5	RO	0	<p>VGAPSE: VGA palette snoop Enable This host bridge does not implement this bit. This bit is hardwired to 0.</p>
4	RO	0	<p>MWIEN: Memory Write and Invalidate Enable This host bridge will never issue memory write and invalidate commands. This bit is therefore hardwired to 0.</p>
3	RO	0	<p>SCE: Special Cycle Enable This host bridge does not implement this bit. This bit is hardwired to a 0.</p>
2	RO	1	<p>BME: Bus Master Enable This host bridge is always enabled as a master. This bit is hardwired to a 1.</p>
1	RO	1	<p>MSE: Memory Space Enable This host bridge always allows access to main memory. This bit is not implemented and is hardwired to 1.</p>
0	RO	0	<p>IOAE: Access Enable This bit is not implemented in this host bridge and is hardwired to 0.</p>

2.5.8 PCISTS - PCI Status Register

The PCI Status register is a 16-bit status register that reports the occurrence of various error events on this device's PCI interface.

Device: 0	Function: 0-1	Offset: 06h	
Device: 2	Function: 0-1, 4-5	Offset: 06h	
Device: 3	Function: 0-2, 4	Offset: 06h	
Device: 4-6	Function: 0-3	Offset: 06h	
Bit	Type	Reset Value	Description
15	RO	0	Detect Parity Error (DPE) The host bridge does not implement this bit and is hardwired to a 0.
14	RO	0	Signaled System Error (SSE) This bit is set to 1 when this device generates an SERR message over the bus for any enabled error condition. If the host bridge does not signal errors using this bit, this bit is hardwired to a 0 and is read only.
13	RO	0	Received Master Abort Status (RMAS) This bit is set when this device generates request that receives an Unsupported Request completion packet. Software clears the bit by writing 1 to it. If this device does not receive Unsupported Request completion packets, the bit is hardwired to 0 and is read only.
12	RO	0	Received Target Abort Status (RTAS) This bit is set when this device generates a request that receives a Completer Abort completion packet. Software clears this bit by writing a 1 to it. If this device does not receive Completer Abort completion packets, this bit is hardwired to 0 and read only.
11	RO	0	Signaled Target Abort Status (STAS) This device will not generate a Target Abort completion or Special Cycle. This bit is not implemented in this device and is hardwired to a 0.
10:9	RO	0	DEVSEL Timing (DEVT) These bits are hardwired to 00. This device does not physically connect to PCI bus X. These bits are set to "00" (fast decode) so that optimum DEVSEL timing for PCI bus X is not limited by this device.
8	RO	0	Master Data Parity Error Detected (DPD) PERR signaling and messaging are not implemented by this bridge, therefore this bit is hardwired to 0.
7	RO	1	Fast Back-to-Back (FB2B) This bit is hardwired to 1. This device is not physically connected to a PCI bus. This bit is set to 1 (indicating back-to-back capabilities) so that the optimum setting for this PCI bus is not limited by this device.
6	RO	0	Reserved
5	RO	0	66 MHz Capable Does not apply to PCI Express. Hardwired to 0.

Device: 0	Function: 0-1	Offset: 06h	
Device: 2	Function: 0-1, 4-5	Offset: 06h	
Device: 3	Function: 0-2, 4	Offset: 06h	
Device: 4-6	Function: 0-3	Offset: 06h	
Bit Type Reset Value Description			
4	RO	TBD	<p>Capability List (CLIST)</p> <p>This bit is hardwired to 1 to indicate to the configuration software that this device/function implements a list of new capabilities. A list of new capabilities is accessed via registers CAPPTR at the configuration address offset 34h from the start of the PCI configuration space header of this function. Register CAPPTR contains the offset pointing to the start address with configuration space of this device where the capability register resides. This bit must be set for a PCI Express device or if the VSEC capability.</p> <p>If no capability structures are implemented, this bit is hardwired to 0.</p>
3	RO	0	<p>Interrupt Status</p> <p>If this device generates an interrupt, then this read-only bit reflects the state of the interrupt in the device/function. Only when the Interrupt Disable bit in the command register is a 0 and this Interrupt Status bit is a 1, will the device's/function's INTx# signal be asserted. Setting the Interrupt Disable bit to a 1 has no effect on the state of this bit.</p> <p>If this device does not generate interrupts, then this bit is not implemented (RO and reads returns 0).</p>
2:0	RO	0	Reserved

2.6 SAD - System Address Decoder Registers

2.6.1 SAD_PAM0123

This register is for legacy device 0, function 0 at 90h-93h address space.

Device: 0	Function: 1	Offset: 40h	Access as a Dword
Bit Type Reset Value Description			
29:28	RW	0	<p>PAM3_HIENABLE. 0D4000h-0D7FFFh Attribute (HIENABLE).</p> <p>This field controls the steering of read and write cycles that address the BIOS area from 0D4000h to 0D7FFFh.</p> <p>00 = DRAM Disabled: All accesses are directed to ESI.</p> <p>01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI.</p> <p>10 = Write Only: All writes are sent to DRAM. Reads are serviced by ESI.</p> <p>11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.</p>

Device: 0 Function: 1 Offset: 40h Access as a Dword			
Bit	Type	Reset Value	Description
25:24	RW	0	PAM3_LOENABLE. 0D0000h-0D3FFFh Attribute (LOENABLE). This field controls the steering of read and write cycles that address the BIOS area from 0D0000h to 0D3FFFh. 00 = DRAM Disabled: All accesses are directed to ESI. 01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI. 10 = Write Only: All writes are send to DRAM. Reads are serviced by ESI. 11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.
21:20	RW	0	PAM2_HIENABLE. 0CC000h-0CFFFFh Attribute (HIENABLE). This field controls the steering of read and write cycles that address the BIOS area from 0CC000h to 0CFFFFh. 00 = DRAM Disabled: All accesses are directed to ESI. 01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI. 10 = Write Only: All writes are send to DRAM. Reads are serviced by ESI. 11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.
17:16	RW	0	PAM2_LOENABLE. 0C8000h-0CBFFFh Attribute (LOENABLE). This field controls the steering of read and write cycles that address the BIOS area from 0C8000h to 0CBFFFh. 00 = DRAM Disabled: All accesses are directed to ESI. 01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI. 10 = Write Only: All writes are send to DRAM. Reads are serviced by ESI. 11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.
13:12	RW	0	PAM1_HIENABLE. 0C4000h-0C7FFFh Attribute (HIENABLE). This field controls the steering of read and write cycles that address the BIOS area from 0C4000h to 0C7FFFh. 00 = DRAM Disabled: All accesses are directed to ESI. 01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI. 10 = Write Only: All writes are send to DRAM. Reads are serviced by ESI. 11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.
9:8	RW	0	PAM1_LOENABLE. 0C0000h-0C3FFFh Attribute (LOENABLE). This field controls the steering of read and write cycles that address the BIOS area from 0C0000h to 0C3FFFh. 00 = DRAM Disabled: All accesses are directed to ESI. 01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI. 10 = Write Only: All writes are send to DRAM. Reads are serviced by ESI. 11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.
5:4	RW	0	PAM0_HIENABLE. 0F0000h-0FFFFFh Attribute (HIENABLE). This field controls the steering of read and write cycles that address the BIOS area from 0F0000h to 0FFFFFh. 00 = DRAM Disabled: All accesses are directed to ESI. 01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI. 10 = Write Only: All writes are send to DRAM. Reads are serviced by ESI. 11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.

2.6.2 SAD_PAM456

Register for legacy device 0, function 0 94h-97h address space.

Device: 0 Function: 1 Offset: 44h Access as a Dword			
Bit	Type	Reset Value	Description
21:20	RW	0	PAM6_HIENABLE. 0EC000h-0EFFFFh Attribute (HIENABLE). This field controls the steering of read and write cycles that address the BIOS area from 0EC000h to 0EFFFFh. 00 = DRAM Disabled: All accesses are directed to ESI. 01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI. 10 = Write Only: All writes are send to DRAM. Reads are serviced by ESI. 11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.
17:16	RW	0	PAM6_LOENABLE. 0E8000-0EBFFF Attribute (LOENABLE). This field controls the steering of read and write cycles that address the BIOS area from 0E8000 to 0EBFFF. 00 = DRAM Disabled: All accesses are directed to ESI. 01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI. 10 = Write Only: All writes are send to DRAM. Reads are serviced by ESI. 11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.
13:12	RW	0	PAM5_HIENABLE. 0E4000h-0E7FFFh Attribute (HIENABLE). This field controls the steering of read and write cycles that address the BIOS area from 0E4000h to 0E7FFFh. 00 = DRAM Disabled: All accesses are directed to ESI. 01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI. 10 = Write Only: All writes are send to DRAM. Reads are serviced by ESI. 11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.
9:8	RW	0	PAM5_LOENABLE. 0E0000h-0E3FFFh Attribute (LOENABLE). This field controls the steering of read and write cycles that address the BIOS area from 0E0000h to 0E3FFFh. 00 = DRAM Disabled: All accesses are directed to ESI. 01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI. 10 = Write Only: All writes are send to DRAM. Reads are serviced by ESI. 11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.
5:4	RW	0	PAM4_HIENABLE. 0DC000h-0DFFFFh Attribute (HIENABLE). This field controls the steering of read and write cycles that address the BIOS area from 0DC000h to 0DFFFFh. 00 = DRAM Disabled: All accesses are directed to ESI. 01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI. 10 = Write Only: All writes are send to DRAM. Reads are serviced by ESI. 11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.
1:0	RW	0	PAM4_LOENABLE. 0D8000h-0DBFFFh Attribute (LOENABLE). This field controls the steering of read and write cycles that address the BIOS area from 0D8000h to 0DBFFFh. 00 = DRAM Disabled: All accesses are directed to ESI. 01 = Read Only: All reads are sent to DRAM. All writes are forwarded to ESI. 10 = Write Only: All writes are send to DRAM. Reads are serviced by ESI. 11 = Normal DRAM Operation: All reads and writes are serviced by DRAM.

2.6.3 SAD_HEN

Register for legacy Hole Enable.

Device: 0 Function: 1 Offset: 48h Access as a Dword			
Bit	Type	Reset Value	Description
7	RW	0	HEN: Hole Enable This field enables a memory hole in DRAM space. The DRAM that lies "behind" this space is not remapped. 0 = No Memory hole. 1 = Memory hole from 15 MB to 16 MB.

2.6.4 SAD_SMRAM

Register for legacy 9Dh address space. Note both IOH and non-core have this now.

Device: 0 Function: 1 Offset: 4Ch Access as a Dword			
Bit	Type	Reset Value	Description
14	RW	0	SMM Space Open (D_OPEN) When D_OPEN=1 and D_LCK=0, the SMM space DRAM is made visible even when SMM decode is not active. This is intended to help BIOS initialize SMM space. Software should ensure that D_OPEN=1 and D_CLS=1 are not set at the same time.
13	RW	0	SMM Space Closed (D_CLS) When D_CLS = 1 SMM space DRAM is not accessible to data references, even if SMM decode is active. Code references may still access SMM space DRAM. This will allow SMM software to reference through SMM space to update the display even when SMM is mapped over the VGA range. Software should ensure that D_OPEN=1 and D_CLS=1 are not set at the same time.
12	RW1S	0	SMM Space Locked (D_LCK) When D_LCK is set to 1 then D_OPEN is reset to 0 and D_LCK, D_OPEN, C_BASE_SEG, G_SMRAME, PCIEBAR, (DRAM_RULEs and INTERLEAVE_LISTs) become read only. D_LCK can be set to 1 via a normal configuration space write but can only be cleared by a Reset. The combination of D_LCK and D_OPEN provide convenience with security. The BIOS can use the D_OPEN function to initialize SMM space and then use D_LCK to "lock down" SMM space in the future so that no application software (or BIOS itself) can violate the integrity of SMM space, even if the program has knowledge of the D_OPEN function. Note that TAD does not implement this lock.
11	RW	0	Global SDRAM Enable (G_SMRAME) If set to a 1, then Compatible SDRAM functions are enabled, providing 128 KB of DRAM accessible at the A0000h address while in SMM (ADSB with SMM decode). To enable Extended SDRAM function this bit has to be set to 1. Once D_LCK is set, this bit becomes read only.
10:8	RO	-	Compatible SMM Space Base Segment (C_BASE_SEG) This field indicates the location of SMM space. SMM DRAM is not remapped. It is simply made visible if the conditions are right to access SMM space, otherwise the access is forwarded to HI. Only SMM space between A0000h and BFFFFh is supported so this field is hardwired to 010.

2.6.5 SAD_PCIEBAR

Global register for PCIEBAR address space.

Device: 0 Function: 1 Offset: 50h Access as a Qword			
Bit	Type	Reset Value	Description
39:20	RW	0	ADDRESS. Base address of PCIEBAR. Must be naturally aligned to size; low order bits are ignored.
3:1	RW	0	SIZE. Size of the PCIEBAR address space. (MAX bus number). 000 = 256 MB. 001 = Reserved. 010 = Reserved. 011 = Reserved. 100 = Reserved. 101 = Reserved. 110 = 64 MB. 111 = 128 MB.
0	RW	0	ENABLE. Enable for PCIEBAR address space. Editing size should not be done without also enabling range.

2.6.6 SAD_DRAM_RULE_0, SAD_DRAM_RULE_1, SAD_DRAM_RULE_2, SAD_DRAM_RULE_3, SAD_DRAM_RULE_4, SAD_DRAM_RULE_5, SAD_DRAM_RULE_6, SAD_DRAM_RULE_7

This register provides SAD DRAM rules. Address Map for package determination.

Device: 0 Function: 1 Offset: 80h, 84h, 88h, 8Ch, 90h, 94h, 98h, 9Ch Access as a Dword			
Bit	Type	Reset Value	Description
19:6	RW	-	LIMIT DRAM rule top limit address. Must be strictly greater than previous rule, even if this rule is disabled, unless this rule and all following rules are disabled. Lower limit is the previous rule (or 0 if it is first rule). This field is compared against MA[39:26] in the memory address map.
2:1	RW	-	MODE DRAM rule interleave mode. If a DRAM_RULE hits a 3 bit number is used to index into the corresponding interleave_list to determine which package the DRAM belongs to. This mode selects how that number is computed. 00 = Address bits {8,7,6}. 01 = Address bits {8,7,6} XORed with {18,17,16}. 10 = Address bit {6}, MOD3(Address[39..6]). (Note 6 is the high order bit) 11 = Reserved.
0	RW	0	ENABLE Enable for DRAM rule. If Enabled Range between this rule and previous rule is Directed to HOME channel (unless overridden by other dedicated address range registers). If disabled, all accesses in this range are directed in MMIO to the IOH.

2.6.7 SAD_INTERLEAVE_LIST_0, SAD_INTERLEAVE_LIST_1, SAD_INTERLEAVE_LIST_2, SAD_INTERLEAVE_LIST_3, SAD_INTERLEAVE_LIST_4, SAD_INTERLEAVE_LIST_5, SAD_INTERLEAVE_LIST_6, SAD_INTERLEAVE_LIST_7

This register provides SAD DRAM package assignments. When the corresponding DRAM_RULE hits, a 3-bit number (determined by mode) is used to index into the interleave_list to determine which package is the HOME for this address.

- 00: IOH
- 01: Socket 0
- 10: Socket 1
- 11: Reserved

Device: 0 Function: 1 Offset: C0h, C4h, C8h, CCh, D0h, D4h, D8h, DCh Access as a Dword			
Bit	Type	Reset Value	Description
29:28	RW	-	PACKAGE7 . Package for index value 7 of interleaves.
25:24	RW	-	PACKAGE6 . Package for index value 6 of interleaves.
21:20	RW	-	PACKAGE5 . Package for index value 5 of interleaves.

Device: 0 Function: 1 Offset: C0h, C4h, C8h, CCh, D0h, D4h, D8h, DCh Access as a Dword			
17:16	RW	-	PACKAGE4. Package for index value 4 of interleaves.
13:12	RW	-	PACKAGE3. Package for index value 3 of interleaves.
9:8	RW	-	PACKAGE2. Package for index value 2 of interleaves.
5:4	RW	-	PACKAGE1. Package for index value 1 of interleaves.
1:0	RW	-	PACKAGE0. Package for index value 0 of interleaves.

2.7 Intel® QuickPath Interconnect Link Registers

2.7.1 QPI_QPILCL_L0, QPI_QPILCL_L1

This register provides Intel QPI Link Control.

Device: 2 Function: 0, 4 Offset: 48h Access as a Dword			
Bit	Type	Reset Value	Description
21	RW	0	L1_MASTER Indicates that this end of the link is the L1 master. This link transmitter bit is an L1 power state master and can initiate an L1 power state transition. If this bit is not set, then the link transmitter is an L1 power state slave and should respond to L1 transitions with an ACK or NACK. If the link power state of L1 is enabled, then there is one master and one slave per link. The master may only issue single L1 requests, while the slave can only issue single L1_Ack or L1_NAck responses for the corresponding request.
20	RW	0	L1_ENABLE Enables L1 mode at the transmitter. This bit should be ANDed with the receive L1 capability bit received during parameter exchange to determine if a transmitter is allowed to enter into L1. This is NOT a bit that determines the capability of a device.
18	RW	0	LOS_ENABLE Enables LOS mode at the transmitter. This bit should be ANDed with the receive LOS capability bit received during parameter exchange to determine if a transmitter is allowed to enter into LOS. This is NOT a bit that determines the capability of a device.

2.8 Integrated Memory Controller Control Registers

The registers in this section apply only to processors supporting registered DIMMs.

2.8.1 MC_CONTROL

This register is the Primary control register.

Device: 3 Function: 0 Offset: 48h Access as a Dword			
Bit	Type	Reset Value	Description
10	RW	0	CHANNEL2_ACTIVE When set, indicates MC channel 2 is active. This bit is controlled (set/reset) by software only. This bit is required to be set for any active channel when INIT_DONE is set by software.
9	RW	0	CHANNEL1_ACTIVE When set, indicates MC channel 1 is active. This bit is controlled (set/reset) by software only. This bit is required to be set for any active channel when INIT_DONE is set by software. Channel 0 AND Channel 1 active must both be set for a lockstep or mirrored pair.
8	RW	0	CHANNEL0_ACTIVE When set, indicate MC channel 0 is active. This bit is controlled (set/reset) by software only. This bit is required to be set for any active channel when INIT_DONE is set by software. Channel 0 AND Channel 1 active must both be set for a lockstep or mirrored pair.
7	WO	0	INIT_DONE MC initialize complete signal. Setting this bit will exit the training mode of the Integrated Memory Controller and begin normal operation including all enabled maintenance operations. Any CHANNEL_ACTIVE bits not set when writing a 1 to INIT_DONE will cause the corresponding channel to be disabled.
6	RW	0	DIVBY3EN Divide By 3 enable. When set, MAD would use the longer pipeline for transactions that are 3 or 6 way interleaved and shorter pipeline for all other transactions. The SAG registers must be appropriately programmed as well.
5	RW	0	CHANNELRESET2 Reset only the state within the channel. Equivalent to pulling warm reset for that channel.
4	RW	0	CHANNELRESET1 Reset only the state within the channel. Equivalent to pulling warm reset for that channel.
3	RW	0	CHANNELRESET0 Reset only the state within the channel. Equivalent to pulling warm reset for that channel.
2	RW	0	AUTOPRECHARGE. Autoprecharge enable. This bit should be set with the closed page bit. If it is not set with closed page, address decode will be done without setting the autoprecharge bit.
1	RW	0	ECCEN: ECC Enable ECC Checking enables. When this bit is set in lockstep mode the ECC checking is for the x8 SDDC. ECCEN without Lockstep enables the x4 SDDC ECC checking.
0	RW	0	CLOSED_PAGE When set, the MC supports a Closed Page policy. The default is Open Page but BIOS should always configure this bit.

2.8.2 MC_STATUS

This register is the MC primary status register.

Device: 3 Function: 0 Offset: 4Ch Access as a Dword			
Bit	Type	Reset Value	Description
4	RO	1	ECC_ENABLED . ECC is enabled.
2	RO	0	CHANNEL2_DISABLED Channel 2 is disabled. This can be factory configured or if Init done is written without the channel_active being set. Clocks in the channel will be disabled when this bit is set.
1	RO	0	CHANNEL1_DISABLED Channel 1 is disabled. This can be factory configured or if Init done is written without the channel_active being set. Clocks in the channel will be disabled when this bit is set.
0	RO	0	CHANNEL0_DISABLED Channel 0 is disabled. This can be factory configured or if Init done is written without the channel_active being set. Clocks in the channel will be disabled when this bit is set.

2.8.3 MC_SMI_SPARE_DIMM_ERROR_STATUS

SMI sparing DIMM error threshold overflow status register. This bit is set when the per-DIMM error counter exceeds the specified threshold. The bit is reset by BIOS.

Device: 3 Function: 0 Offset: 50h Access as a Dword			
Bit	Type	Reset Value	Description
13:12	RWOC	0	REDUNDANCY_LOSS FAILING_DIMM The ID for the failing DIMM when redundancy is lost.
11:0	RWOC	0	DIMM_ERROR_OVERFLOW_STATUS This 12-bit field is the per dimm error overflow status bits. The organization is as follows: If there are three or more DIMMS on the channel: Bit 0 = DIMM 0 Channel 0 Bit 1 = DIMM 1 Channel 0 Bit 2 = DIMM 2 Channel 0 Bit 3 = DIMM 3 Channel 0 Bit 4 = DIMM 0 Channel 1 Bit 5 = DIMM 1 Channel 1 Bit 6 = DIMM 2 Channel 1 Bit 7 = DIMM 3 Channel 1 Bit 8 = DIMM 0 Channel 2 Bit 9 = DIMM 1 Channel 2 Bit 10 = DIMM 2 Channel 2 Bit 11 = DIMM 3 Channel 2 If there are one or two DIMMS on the channel: Bit 0 = DIMM 0, Ranks 0 and 1, Channel 0 Bit 1 = DIMM 0, Ranks 2 and 3, Channel 0 Bit 2 = DIMM 1, Ranks 0 and 1, Channel 0 Bit 3 = DIMM 1, Ranks 2 and 3, Channel 0 Bit 4 = DIMM 0, Ranks 0 and 1, Channel 1 Bit 5 = DIMM 0, Ranks 2 and 3, Channel 1 Bit 6 = DIMM 1, Ranks 0 and 1, Channel 1 Bit 7 = DIMM 1, Ranks 2 and 3, Channel 1 Bit 8 = DIMM 0, Ranks 0 and 1, Channel 2 Bit 9 = DIMM 0, Ranks 2 and 3, Channel 2 Bit 10 = DIMM 1, Ranks 0 and 1, Channel 2 Bit 11 = DIMM 1, Ranks 2 and 3, Channel 2

2.8.4 MC_SMI_SPARE_CNTRL

System Management Interrupt and Spare control register.

Device: 3 Function: 0 Offset: 54h Access as a Dword			
Bit	Type	Reset Value	Description
16	RW	0	INTERRUPT_SELECT_NMI 1 = Enable NMI signaling. 0 = Disable NMI signaling. If both NMI and SMI enable bits are set, then only SMI is sent.
15	RW	0	INTERRUPT_SELECT_SMI 1 = Enable SMI signaling. 0 = Disable SMI signaling. If both NMI and SMI enable bits are set, then only SMI is sent. This bit functions the same way in Mirror and Independent Modes. The possible SMI events enabled by this bit are: Any one of the error counters MC_COR_ECC_CNT_X meets the value of SMI_ERROR_THRESHOLD field of this register. MC_SSRSTATUS.CMPLT bit is set to 1. MC_RAS_STATUS.REDUNDANCY_LOSS bit is set to 1.
14:0	RW	0	SMI_ERROR_THRESHOLD Defines the error threshold to compare against the per-DIMM error counters MC_COR_ECC_CNT_X, which are also 15 bits.

2.8.5 MC_RESET_CONTROL

DIMM Reset enabling controls.

Device: 3 Function: 0 Offset: 5Ch Access as a Dword			
Bit	Type	Reset Value	Description
0	WO	0	BIOS_RESET_ENABLE When set, MC takes over control of driving RESET to the DIMMs. This bit is set on S3 exit and cold boot to take over RESET driving responsibility from the physical layer.

2.8.6 MC_CHANNEL_MAPPER

Channel mapping register. The sequence of operations to update this register is:

Read MC_Channel_Mapper register

Compare data read to data to be written. If different, then write.

Poll MC_Channel_Mapper register until the data read matches data written.

Device: 3 Function: 0 Offset: 60h Access as a Dword			
Bit	Type	Reset Value	Description
17:15	RW	0	RDLCH2. Mapping of Logical Channel 2 to physical channel for Reads. 001 = Maps to physical Channel 0 010 = Maps to physical Channel 1 100 = Maps to physical Channel 2
14:12	RW	0	WRLCH2. Mapping of Logical Channel 2 to physical channel for Writes. 001 = Maps to physical Channel 0 010 = Maps to physical Channel 1 100 = Maps to physical Channel 2
11:9	RW	0	RDLCH1. Mapping of Logical Channel 1 to physical channel for Reads. 001 = Maps to physical Channel 0 010 = Maps to physical Channel 1 100 = Maps to physical Channel 2
8:6	RW	0	WRLCH1. Mapping of Logical Channel 1 to physical channel for Writes. 001 = Maps to physical Channel 0 010 = Maps to physical Channel 1 100 = Maps to physical Channel 2
5:3	RW	0	RDLCHO. Mapping of Logical Channel 0 to physical channel for Read. 001 = Maps to physical Channel 0 010 = Maps to physical Channel 1 100 = Maps to physical Channel 2
2:0	RW	0	WRLCHO. Mapping of Logical Channel 0 to physical channel for Writes. 001 = Maps to physical Channel 0 010 = Maps to physical Channel 1 100 = Maps to physical Channel 2

2.8.7 MC_MAX_DOD

This register defines the MAX number of DIMMS, RANKS, BANKS, ROWS, COLS among all DIMMS populating the three channels. The Memory Init logic uses this register to cycle through all the memory addresses writing all 0's to initialize all locations. This register is also used for scrubbing and sparing and must always be programmed if any DODs are programmed.

Device: 3 Function: 0 Offset: 64h Access as a Dword			
Bit	Type	Reset Value	Description
10:9	RW	0	MAXNUMCOL. Maximum Number of Columns. 00 = 2^{10} columns 01 = 2^{11} columns 10 = 2^{12} columns 11 = RSVD.
8:6	RW	0	MAXNUMROW. Maximum Number of Rows. 000 = 2^{12} Rows 001 = 2^{13} Rows 010 = 2^{14} Rows 011 = 2^{15} Rows 100 = 2^{16} Rows Others = RSVD.
5:4	RW	0	MAXNUMBANK. Max Number of Banks. 00 = Four-banked 01 = Eight-banked 10 = Sixteen-banked.
3:2	RW	0	MAXNUMRANK. Maximum Number of Ranks. 00 = Single Ranked 01 = Double Ranked 10 = Quad Ranked.
1:0	RW	0	MAXNUMDIMMS. Maximum Number of DIMMs. 00 = 1 DIMM 01 = 2 DIMMs 10 = 3 DIMMs 11 = RSVD.

2.8.8 MC_RD_CRD_T_INIT

These registers contain the initial read credits available for issuing memory reads. TAD read credit counters are loaded with the corresponding values at reset and anytime this register is written. BIOS must initialize this register with appropriate values depending on the level of Isoch support in the platform. It is invalid to write this register while TAD is active (has memory requests outstanding), as the write will break TAD's outstanding credit count values.

Register programming rules:

- Total read credits (CRDT_RD + CRDT_RD_HIGH + CRDT_RD_CRIT) must not exceed 31.
- CRDT_RD_HIGH value must correspond to the number of high RTIDs reserved at the IOH.
- CRDT_RD_CRIT value must correspond to the number of critical RTIDs reserved at the IOH.
- CRDT_RD_HIGH + CRDT_RD must be less than or equal to 13.
- CRDT_RD_HIGH + CRDT_RD_CRIT must be less than or equal to 8.
- CRDT_RD_CRIT must be less than or equal to 6. Set CRDT_RD to (16 - CRDT_RD_CRIT - CRDT_RD_HIGH).
- If (Mirroring OR Sparing enabled) then Max for CRDT_RD is 14, otherwise it is 15.
- If (Isoch not enabled) then CRDT_RD_HIGH and CRDT_RD_CRIT are set to 0.

Device: 3 Function: 0 Offset: 70h Access as a Dword			
Bit	Type	Reset Value	Description
20:16	RW	3	CRDT_RD_CRIT. Critical Read Credits.
12:8	RW	1	CRDT_RD_HIGH. High Read Credits.
4:0	RW	13	CRDT_RD. Normal Read Credits.



2.8.9 MC_CREDT_WR_THLD

This is the Memory Controller Write Credit Thresholds register. A Write threshold is defined as the number of credits reserved for this priority (or higher) request. It is required that High threshold be greater than or equal to Crit threshold, and that both be lower than the total Write Credit init value. BIOS must initialize this register with appropriate values depending on the level of Isoch support in the platform. The new values take effect immediately upon being written.

Register programming rules:

- CRIT threshold value must correspond to the number of critical RTIDs reserved at the IOH.
- HIGH threshold value must correspond to the sum of critical and high RTIDs reserved at the IOH (which must not exceed 30).
- Set MC_Channel_*_WAQ_PARAMS.ISOCENTRYTHRESHOLD equal to (31-CRIT).

Device: 3 Function: 0 Offset: 74h Access as a Dword			
Bit	Type	Reset Value	Description
12:8	RW	4	HIGH. High Credit Threshold.
4:0	RW	3	CRIT. Critical Credit Threshold.

2.8.10 MC_SCRUBADDR_LO

This register contains part of the address of the last patrol scrub request issued. When running Memtest, the failing address is logged in this register on Memtest errors. Software can write the next address to be scrubbed into this register. Patrol scrubs must be disabled to reliably write this register.

Device: 3 Function: 0 Offset: 78h Access as a Dword			
Bit	Type	Reset Value	Description
29:14	RW	0	PAGE. This field contains the row of the last scrub issued. Can be written to specify the next scrub address with STARTSCRUB in the MC_SCRUB_CONTROL register.
13:0	RW	0	COLUMN. This field contains the column of the last scrub issued. Can be written to specify the next scrub address with STARTSCRUB in the MC_SCRUB_CONTROL register.

2.8.11 MC_SCRUBADDR_HI

This register pair contains part of the address of the last patrol scrub request issued. When running memtest, the failing address is logged in this register on memtest errors. Software can write the next address into this register. Scrubbing must be disabled to reliably read and write this register.

Device: 3 Function: 0 Offset: 7Ch Access as a Dword			
Bit	Type	Reset Value	Description
9:8	RW	0	CHNL. This field can be written to specify the next scrub address with STARTSCRUB in the MC_SCRUB_CONTROL register. This register is not updated with channel address of the last scrub address issued.
7:6	RW	0	DIMM. This field contains the DIMM of the last scrub issued. Can be written to specify the next scrub address with STARTSCRUB in the MC_SCRUB_CONTROL register. For writes, to the register this field always contains the Rank ID. For reads, the following translation must be done: If 3 DIMMs are on the channel, then the rank is RANK[0] while the dimm is the concatenation of DIMM[0] and RANK[1].
5:4	RW	0	RANK. This field contains the rank of the last scrub issued. Can be written to specify the next scrub address with STARTSCRUB in the MC_SCRUB_CONTROL register. For writes, to the register this field always contains the rank id. For reads, the following translation must be done: If 3 dimms are on the channel then the rank is RANK[0] while the dimm is the concatenation of DIMM[0] and RANK[1].
3:0	RW	0	BANK. This field contains the bank of the last scrub issued. Can be written to specify the next scrub address with STARTSCRUB in the MC_SCRUB_CONTROL register.

2.9 TAD – Target Address Decoder Registers

2.9.1 **TAD_DRAM_RULE_0, TAD_DRAM_RULE_1** **TAD_DRAM_RULE_2, TAD_DRAM_RULE_3** **TAD_DRAM_RULE_4, TAD_DRAM_RULE_5** **TAD_DRAM_RULE_6, TAD_DRAM_RULE_7**

TAD DRAM rules. Address map for channel determination within a package. All addresses sent to this HOME agent must hit a valid enabled DRAM_RULE. No error will be generated if they do not hit a valid location and memory aliasing will happen.

Device: 3 Function: 1 Offset: 80h, 84h, 88h, 8Ch, 90h, 94h, 98h, 9Ch Access as a Dword			
Bit	Type	Reset Value	Description
19:6	RW	-	LIMIT. DRAM rule top limit address. Must be strictly greater than previous rule, even if this rule is disabled, unless this rule and all following rules are disabled. Lower limit is the previous rule (or 0 if it is the first rule).
2:1	RW	-	MODE. DRAM rule interleave mode. If a DRAM_RULE hits, a 3-bit number is used to index into the corresponding interleave_list to determine which channel the DRAM belongs to. This mode selects how that number is computed. 00 = Address bits {8,7,6}. 01 = Address bits {8,7,6} XORed with {18,17,16}. 10 = Address bit {6}, MOD3(Address[39..6]). (Note 6 is the high order bit) 11 = Reserved.
0	RW	0	ENABLE. Enable for DRAM rule.

2.9.2

**TAD_INTERLEAVE_LIST_0, TAD_INTERLEAVE_LIST_1
 TAD_INTERLEAVE_LIST_2, TAD_INTERLEAVE_LIST_3
 TAD_INTERLEAVE_LIST_4, TAD_INTERLEAVE_LIST_5
 TAD_INTERLEAVE_LIST_6, TAD_INTERLEAVE_LIST_7**

TAD DRAM package assignments. When the corresponding DRAM_RULE hits, a 3-bit number (determined by mode) is used to index into the Interleave_List Branches to determine which channel the DRAM request belongs to.

Device: 3 Function: 1 Offset: C0h, C4h, C8h, CCh, D0h, D4h, D8h, DCh Access as a Dword			
Bit	Type	Reset Value	Description
29:28	RW	-	Logical Channel7. Index 111 of the Interleave List. Bits determined from the matching TAD_DRAM_RULE mode. 00 = Logical channel 0 01 = Logical channel 1 10 = Logical channel 2 11 = Reserved
25:24	RW	-	Logical Channel6. Index 110 of the Interleave List. Bits determined from the matching TAD_DRAM_RULE mode. 00 = Logical channel 0 01 = Logical channel 1 10 = Logical channel 2 11 = Reserved
21:20	RW	-	Logical Channel5. Index 101 of the Interleave List. Bits determined from the matching TAD_DRAM_RULE mode. 00 = Logical channel 0 01 = Logical channel 1 10 = Logical channel 2 11 = Reserved
17:16	RW	-	Logical Channel4. Index 100 of the Interleave List. Bits determined from the matching TAD_DRAM_RULE mode. 00 = Logical channel 0 01 = Logical channel 1 10 = Logical channel 2 11 = Reserved
13:12	RW	-	Logical Channel3. Index 011 of the Interleave List. Bits determined from the matching TAD_DRAM_RULE mode. 00 = Logical channel 0 01 = Logical channel 1 10 = Logical channel 2 11 = Reserved

Device: 3 Function: 1 Offset: C0h, C4h, C8h, CCh, D0h, D4h, D8h, DCh Access as a Dword			
Bit	Type	Reset Value	Description
9:8	RW	-	Logical Channel2. Index 010 of the Interleave List. Bits determined from the matching TAD_DRAM_RULE mode. 00 = Logical channel 0 01 = Logical channel 1 10 = Logical channel 2 11 = Reserved
5:4	RW	-	Logical Channel1. Index 001 of the Interleave List. Bits determined from the matching TAD_DRAM_RULE mode. 00 = Logical channel 0 01 = Logical channel 1 10 = Logical channel 2 11 = Reserved
1:0	RW	-	Logical Channel0. Index 000 of the Interleave List. Bits determined from the matching TAD_DRAM_RULE mode. 00 = Logical channel 0 01 = Logical channel 1 10 = Logical channel 2 11 = Reserved

2.10 Integrated Memory Controller Channel Control Registers

2.10.1 MC_CHANNEL_0_DIMM_RESET_CMD

Integrated Memory Controller DIMM reset command register. This register is used to sequence the reset signals to the DIMMs.

Device: 4, 5, 6 Function: 0 Offset: 50h Access as a Dword			
Bit	Type	Reset Value	Description
2	RW	0	BLOCK_CKE. When set, CKE will be forced to be deasserted.
1	RW	0	ASSERT_RESET. When set, Reset will be driven to the DIMMs.
0	WO	0	RESET. Reset the DIMMs. Setting this bit will cause the Integrated Memory Controller DIMM Reset state machine to sequence through the reset sequence using the parameters in MC_DIMM_INIT_PARAMS.

2.10.2 MC_CHANNEL_0_DIMM_INIT_CMD MC_CHANNEL_1_DIMM_INIT_CMD MC_CHANNEL_2_DIMM_INIT_CMD

Integrated Memory Controller DIMM initialization command register. This register is used to sequence the channel through the physical layer training required for DDR.

Device: 4, 5, 6 Function: 0 Offset: 54h Access as a Dword			
Bit	Type	Reset Value	Description
17	WO	0	ASSERT_CKE. When set, all CKE will be asserted. Write a 0 to this bit to stop the init block from driving CKE. This bit has no effect once MC_CONTROL.INIT_DONE is set. This bit must be used during INITIALIZATION only and be cleared out before MC_CONTROL.INIT_DONE is set. This bit must not be asserted during initialization for S3 resume.
16	RW	0	DO_RCOMP. When set, an RCOMP will be issued to the rank specified in the RANK field.
15	RW	0	DO_ZQCL. When set, a ZQCL will be issued to the rank specified in the RANK field.
14	RW	0	WRDQDQS_MASK. When set, the Write DQ-DQS training will be skipped.
13	RW	0	WRLEVEL_MASK. When set, the Write Levelization step will be skipped.
12	RW	0	RDDQDQS_MASK. When set, the Read DQ-DQS step will be skipped.
11	RW	0	RCVEN_MASK. When set, the RCVEN step will be skipped.
10	WO	0	RESET_FIFOS. When set, the TX and RX FIFO pointers will be reset at the next BCLK edge. The Bubble Generators will also be reset.
9	RW	0	IGNORE_RX. When set, the read return datapath will ignore all data coming from the RX FIFOs. This is done by gating the early valid bit.
8	RW	0	STOP_ON_FAIL. When set along with the AUTORESETDIS not being set, the phyinit FSM will stop if a step has not completed after timing out.
7:5	RW	0	RANK. The rank currently being tested. The PhyInit FSM must be sequenced for every rank present in the channel. The rank value is set to the rank being trained.
4:2	RW	0	NXT_PHYINIT_STATE. Set to sequence the physical layer state machine. 000 = IDLE 001 = RD DQ-DQS 010 = RcvEn Bitlock 011 = Write Level 100 = WR DQ-DQS.
1	RW	0	AUTODIS. Disables the automatic training where each step is automatically incremented. When set, the physical layer state machine must be sequenced with software. The training FSM must be sequenced using the NXT_PHYINIT_STATE field.
0	WO	0	TRAIN. Cycle through the training sequence for the rank specified in the RANK field.

2.10.3 MC_CHANNEL_0_DIMM_INIT_PARAMS MC_CHANNEL_1_DIMM_INIT_PARAMS MC_CHANNEL_2_DIMM_INIT_PARAMS

Initialization sequence parameters are stored in this register. Each field is 2^n count.

Device: 4, 5, 6 Function: 0 Offset: 58h Access as a Dword			
Bit	Type	Reset Value	Description
26	RW	0	DIS_3T. When set, 3T mode will not be enabled as a part of the MRS write to the RDIMM. The RC2 write to switch to 3T and back to 1T timing before and after an MRS write will not be done if the bit is set. This bit should be set if the RDIMM supports auto MRS cycles where the dimm takes care of the 3T switching on MRS writes.
25	RW	0	DIS_AI. When set, address inversion will not be disabled as a part of the MRS write to the RDIMM. The RCO write to disable and enable address inversion will not be done. This bit should be set if the RDIMM supports auto MRS cycles where the dimm takes care of disabling address inversion for MRS writes.
24	RW	0	THREE_DIMMS_PRESENT. Set when channel contains three DIMMs. THREE_DIMMS_PRESENT=1 and QUAD_RANK_PRESENT=1 (or SINGLE_QUAD_RANK_PRESENT=1) are mutually exclusive.
23	RW	0	SINGLE_QUAD_RANK_PRESENT. Set when channel contains a single quad rank DIMM.
22	RW	0	QUAD_RANK_PRESENT. Set when channel contains 1 or 2 quad rank DIMMs.
21:17	RW	15	WRDQDQS_DELAY. Specifies the delay in DCLKs between reads and writes for WRDQDQS training.
16	RW	0	WRLEVEL_DELAY. Specifies the delay used between write CAS indications for write leveling training. 0 = 16 DCLKs. 1 = 32 DCLKs.
15	RW	0	REGISTERED_DIMM. Set when channel contains registered DIMMs.
14:10	RW	0	PHY_FSM_DELAY. Global timer used for bounding the physical layer training. If the timer expires, the FSM will go to the next step and the counter will be reloaded with PHY_FSM_DELAY value. Units are 2^n dclk.
9:5	RW	0	BLOCK_CKE_DELAY. Delay in ns from when clocks and command are valid to the point CKE is allowed to be asserted. Units are in 2^n uclk.
4:0	RW	0	RESET_ON_TIME. Reset will be asserted for the time specified. Units are 2^n Uclk.

2.10.4 MC_CHANNEL_0_DIMM_INIT_STATUS MC_CHANNEL_1_DIMM_INIT_STATUS MC_CHANNEL_2_DIMM_INIT_STATUS

The initialization state is stored in this register. This register is cleared on a new training command.

Device: 4, 5, 6 Function: 0 Offset: 5Ch Access as a Dword			
Bit	Type	Reset Value	Description
9	RO	0	RCOMP_CMPLT. When set, indicates that RCOMP command has complete. This bit is cleared by hardware on command issuance and set once the command is complete.
8	RO	0	INIT_CMPLT. This bit is cleared when a new training command is issued. It is set once the sequence is complete regardless of whether all steps passed or not.
7	RO	0	ZQCL_CMPLT. When set, indicates that ZQCL command has completed. This bit is cleared by hardware on command issuance and set once the command is complete.
6	RO	0	WR_DQ_DQS_PASS. Set after a training command when the Write DQ-DQS training step passes. The bit is cleared by hardware when a new training command is sent.
5	RO	0	WR_LEVEL_PASS. Set after a training command when the write leveling training step passes. The bit is cleared by hardware when a new training command is sent.
4	RO	0	RD_RCVEN_PASS. Set after a training command when the Read Receive Enable training step passes. The bit is cleared by hardware when a new training command is sent.
3	RO	0	RD_DQ_DQS_PASS. Set after a training command when the Read DQ-DQS training step passes. The bit is cleared by hardware when a new training command is sent.
2:0	RO	0	PHYFSMSTATE. The current state of the top level training FSM. 000 = IDLE 001 = RD DQ-DQS 010 = RcvEn Bitlock 011 = Write Level 100 = WR DQ-DQS

2.10.5 MC_CHANNEL_0_DDR3CMD MC_CHANNEL_1_DDR3CMD MC_CHANNEL_2_DDR3CMD

DDR3 Configuration Command. This register is used to issue commands to the DIMMs such as MRS commands. The register is used by setting one of the *_VALID bits along with the appropriate address and destination RANK. The command is then issued directly to the DIMM. Care must be taken in using this register as there is no enforcement of timing parameters related to the action taken by a DDR3CMD write. This register has no effect after MC_CONTROL.INIT_DONE is set.

Device: 4, 5, 6 Function: 0 Offset: 60h Access as a Dword			
Bit	Type	Reset Value	Description
28	RW	0	PRECHARGE_VALID. Indicates current command is for a precharge command.
27	RW	0	ACTIVATE_VALID. Indicates current command is for an activate command.
26	RW	0	REG_VALID. Indicates current command is for a registered DIMM config write Bit is cleared by hardware on issuance. This bit applies only to processors supporting registered DIMMs.
25	RW	0	WR_VALID. Indicates current command is for a write CAS. Bit is cleared by hardware on issuance.
24	RW	0	RD_VALID. Indicates current command is for a read CAS. Bit is cleared by hardware on issuance.
23	RW	0	MRS_VALID. Indicates current command is an MRS command. Bit is cleared by hardware on issuance.
22:20	RW	0	RANK. Destination rank for command.
19:16	RW	0	MRS_BA. Address bits driven to DDR_BA[2:0] pins for the DRAM command being issued due to a valid bit being set in this register.
15:0	RW	0	MRS_ADDR. Address bits driven to DDR_MA pins for the DRAM command being issued due to a valid bit being set in this register.

2.10.6 MC_CHANNEL_0_REFRESH_THROTTLE_SUPPORT MC_CHANNEL_1_REFRESH_THROTTLE_SUPPORT MC_CHANNEL_2_REFRESH_THROTTLE_SUPPORT

This register supports Self Refresh and Thermal Throttle functions.

Device: 4, 5, 6 Function: 0 Offset: 68h Access as a Dword			
Bit	Type	Reset Value	Description
3:2	RW	0	INC_ENTERPWRDWN_RATE. Powerdown rate will be increased during thermal throttling based on the following configurations. 00 = tRANKIDLE (Default) 01 = 16 10 = 24 11 = 32
1	RW	0	DIS_OP_REFRESH. When set, the refresh engine will not issue opportunistic refresh.
0	RW	0	ASR_PRESENT. When set, indicates DRAMs on this channel can support Automatic Self Refresh. If the DRAM is not supporting ASR (Auto Self Refresh), then Self Refresh entry will be delayed until the temperature is below the 2x refresh temperature.

2.10.7 MC_CHANNEL_0_MRS_VALUE_0_1 MC_CHANNEL_1_MRS_VALUE_0_1 MC_CHANNEL_2_MRS_VALUE_0_1

The initial MRS register values for MRO, and MR1 can be specified in this register. These values are used for the automated MRS writes used as a part of the training FSM. The remaining values of the MRS register must be specified here.

Device: 4, 5, 6 Function: 0 Offset: 70h Access as a Dword			
Bit	Type	Reset Value	Description
31:16	RW	0	MR1. The values to write to MR1 for A15:A0.
15:0	RW	0	MRO. The values to write to MRO for A15:A0.

2.10.8 MC_CHANNEL_0_MRS_VALUE_2 MC_CHANNEL_1_MRS_VALUE_2 MC_CHANNEL_2_MRS_VALUE_2

The initial MRS register values for MR2. This register also contains the values used for RCO and RC2 writes for registered DIMMs. These values are used during the automated training sequence when MRS writes or registered DIMM RC writes are used. The RC fields do not need to be programmed if the address inversion and 3T/1T transitions are disabled.

Device: 4, 5, 6 Function: 0 Offset: 74h Access as a Dword			
Bit	Type	Reset Value	Description
23:20	RW	0	RC2. The values to write to the RC2 register on RDIMMs. This value will be written whenever 3T or 1T timings are enabled by hardware. For this reason bit 1 of the RC2 field (bit 21 of this register) will be controlled by hardware. [23:22] and [20] will be driven with the RDIMM register write command for RC2.
19:16	RW	0	RC0. The values to write to the RCO register on RDIMMs. This value will be written whenever address inversion is enabled or disabled by hardware. For this reason bit 0 of the RCO field (bit 16 of this register) will be controlled by hardware. [19:17] will be driven with the RDIMM register write command for RCO.
15:0	RW	0	MR2. The values to write to MR2 for A15:A0.

2.10.9 MC_CHANNEL_0_RANK_PRESENT MC_CHANNEL_1_RANK_PRESENT MC_CHANNEL_2_RANK_PRESENT

This register provides the rank present vector.

Device: 4, 5, 6 Function: 0 Offset: 7Ch Access as a Dword			
Bit	Type	Reset Value	Description
7:0	RW	0	RANK_PRESENT. Vector that represents the ranks that are present. Each bit represents a logical rank. When two or fewer DIMMs are present, [3:0] represents the four possible ranks in DIMM0 and [7:4] represents the ranks that are possible in DIMM1. When three DIMMs are present, then the following applies: [1:0] represents ranks 1:0 in Slot 0 [3:2] represents ranks 3:2 in Slot 1 [5:4] represents ranks 5:4 in Slot 2

2.10.10 MC_CHANNEL_0_RANK_TIMING_A MC_CHANNEL_1_RANK_TIMING_A MC_CHANNEL_2_RANK_TIMING_A

This register contains parameters that specify the rank timing used. All parameters are in DCLK.

Device: 4, 5, 6 Function: 0 Offset: 80h Access as a Dword			
Bit	Type	Reset Value	Description
28:26	RW	0	tddWrTRd. Minimum delay between a write followed by a read to different DIMMs. 000 = 1 001 = 2 010 = 3 011 = 4 100 = 5 101 = 6 110 = 7 111 = 8
25:23	RW	0	tdrWrTRd. Minimum delay between a write followed by a read to different ranks on the same DIMM. 000 = 1 001 = 2 010 = 3 011 = 4 100 = 5 101 = 6 110 = 7 111 = 8
22:19	RW	0	tsrWrTRd. Minimum delay between a write followed by a read to the same rank. 0000 = 10 0001 = 11 0010 = 12 0011 = 13 0100 = 14 0101 = 15 0110 = 16 0111 = 17 1000 = 18 1001 = 19 1010 = 20 1011 = 21 1100 = 22

Device: 4, 5, 6 Function: 0 Offset: 80h Access as a Dword			
Bit	Type	Reset Value	Description
18:15	RW	0	tddRdTWr. Minimum delay between Read followed by a Write to different DIMMs. 0000 = 2 0001 = 3 0010 = 4 0011 = 5 0100 = 6 0101 = 7 0110 = 8 0111 = 9 1000 = 10 1001 = 11 1010 = 12 1011 = 13 1100 = 14
14:11	RW	0	tdrRdTWr. Minimum delay between Read followed by a write to different ranks on the same DIMM. 0000 = 2 0001 = 3 0010 = 4 0011 = 5 0100 = 6 0101 = 7 0110 = 8 0111 = 9 1000 = 10 1001 = 11 1010 = 12 1011 = 13 1100 = 14
10:7	RW	0	tsrRdTWr. Minimum delay between Read followed by a write to the same rank. 0000 = RSVD 0001 = RSVD 0010 = RSVD 0011 = 5 0100 = 6 0101 = 7 0110 = 8 0111 = 9 1000 = 10 1001 = 11 1010 = 12 1011 = 13 1100 = 14

Device: 4, 5, 6 Function: 0 Offset: 80h Access as a Dword			
Bit	Type	Reset Value	Description
6:4	RW	0	tddRdTRd. Minimum delay between reads to different DIMMs. 000 = 2 001 = 3 010 = 4 011 = 5 100 = 6 101 = 7 110 = 8 111 = 9
3:1	RW	0	tdrRdTRd. Minimum delay between reads to different ranks on the same DIMM. 000 = 2 001 = 3 010 = 4 011 = 5 100 = 6 101 = 7 110 = 8 111 = 9
0	RW	0	tsrRdTRd. Minimum delay between reads to the same rank. 0 = 4 1 = 6

2.10.11 MC_CHANNEL_0_RANK_TIMING_B MC_CHANNEL_1_RANK_TIMING_B MC_CHANNEL_2_RANK_TIMING_B

This register contains parameters that specify the rank timing used. All parameters are in DCLK.

Device: 4, 5, 6 Function: 0 Offset: 84h Access as a Dword			
Bit	Type	Reset Value	Description
20:16	RW	0	B2B_CAS_DELAY. Controls the delay between CAS commands in DCLKS. The minimum spacing is 4 DCLKS. Values below 3 have no effect. A value of 0 disables the logic. Setting the value between 3-31 also spaces the read data by 0-29 DCLKS. The value entered is one less than the spacing required, i.e. a spacing of 5 DCLKS between CAS commands (or 1 DCLK on the read data) requires a setting of 4.
15:13	RW	0	tddWrTWr. Minimum delay between writes to different DIMMs. 000 = 2 001 = 3 010 = 4 011 = 5 100 = 6 101 = 7 110 = 8 111 = 9
12:10	RW	0	tdrWrTWr. Minimum delay between writes to different ranks on the same DIMM. 000 = 2 001 = 3 010 = 4 011 = 5 100 = 6 101 = 7 110 = 8 111 = 9
9	RW	0	tsrWrTWr. Minimum delay between writes to the same rank. 0 = 4 1 = 6
8:6	RW	0	tRRD. Specifies the minimum time between activate commands to the same rank.
5:0	RW	0	tFAW. Four Activate Window. Specifies the time window in which four activates are allowed the same rank.

2.10.12 MC_CHANNEL_0_BANK_TIMING MC_CHANNEL_1_BANK_TIMING MC_CHANNEL_2_BANK_TIMING

This register contains parameters that specify the bank timing parameters. These values are in DCLK. The values in these registers are encoded where noted. All of these values apply to commands to the same rank only.

Device: 4, 5, 6 Function: 0 Offset: 88h Access as a Dword			
Bit	Type	Reset Value	Description
21:17	RW	0	tWTPr. Minimum Write CAS to Precharge command delay.
16:13	RW	0	tRTPr. Minimum Read CAS to Precharge command delay.
12:9	RW	0	tRCD. Minimum delay between Activate and CAS commands.
8:4	RW	0	tRAS. Minimum delay between Activate and Precharge commands.
3:0	RW	0	tRP. Minimum delay between Precharge command and Activate command.

2.10.13 MC_CHANNEL_0_REFRESH_TIMING MC_CHANNEL_1_REFRESH_TIMING MC_CHANNEL_2_REFRESH_TIMING

This register contains parameters that specify the refresh timings. Units are in DCLK.

Device: 4, 5, 6 Function: 0 Offset: 8Ch Access as a Dword			
Bit	Type	Reset Value	Description
29:19	RW	0	tTHROT_OPREF. The minimum time between two opportunistic refreshes. Should be set to tRFC in DCLKS. Zero is an invalid encoding. A value of 1 should be programmed to disable the throttling of opportunistic refreshes. By setting this field to tRFC, current to a single DIMM can be limited to that required to support this scenario without significant performance impact: <ul style="list-style-type: none"> • 8 panic refreshes in tREFI to one rank • 1 opportunistic refresh every tRFC to another rank • full bandwidth delivered by the third and fourth ranks Platforms that can supply peak currents to the DIMMs should disable opportunistic refresh throttling for max performance.
18:9	RW	0	tREFI_8. Average periodic refresh interval divided by 8.
8:0	RW	0	tRFC. Delay between the refresh command and an activate or refresh command.

2.10.14 MC_CHANNEL_0_CKE_TIMING MC_CHANNEL_1_CKE_TIMING MC_CHANNEL_2_CKE_TIMING

This register contains parameters that specify the CKE timings. All units are in DCLK.

Device: 4, 5, 6 Function: 0 Offset: 90h Access as a Dword			
Bit	Type	Reset Value	Description
31:24	RW	0	tRANKIDLE. Rank will go into powerdown after it has been idle for the specified number of dclks. tRANKIDLE covers max(txXXPDEN). Minimum value is tWRAPDEN. If CKE is being shared between ranks then both ranks must be idle for this amount of time. A Power Down Entry command will be requested for a rank after this number of DCLKs if no request to the rank is in the MC.
23:21	RW	0	tXP. Minimum delay from exit power down with DLL and any valid command. Exit Precharge Power Down with DLL frozen to commands not requiring a locked DLL. Slow exit precharge powerdown is not supported.
20:11	RW	0	tXSDLL. Minimum delay between the exit of self refresh and commands that require a locked DLL.
10:3	RW	0	tXS. Minimum delay between the exit of self refresh and commands not requiring a DLL.
2:0	RW	0	tCKE. CKE minimum pulse width.

2.10.15 MC_CHANNEL_0_ZQ_TIMING MC_CHANNEL_1_ZQ_TIMING MC_CHANNEL_2_ZQ_TIMING

This register contains parameters that specify ZQ timing. All units are DCLK unless otherwise specified. The register encodings are specified where applicable.

Device: 4, 5, 6 Function: 0 Offset: 94h Access as a Dword			
Bit	Type	Reset Value	Description
30	RW	1	Parallel_ZQ. Enable ZQ calibration to different ranks in parallel.
29	RW	1	tZQenable. Enable the issuing of periodic ZQCS calibration commands.
28:8	RW	16410	ZQ_Interval. Nominal interval between periodic ZQ calibration in increments of tREFI.
7:5	RW	4	tZQCS. This field specifies ZQCS cycles in increments of 16. This is the minimum delay between ZQCS and any other command. This register should be programmed to at least 64/16=4='100' to conform to the DDR3 specification.
4:0	RW	0	tZQInit. This field specifies ZQInit cycles in increments of 32. This is the minimum delay between ZQCL and any other command. This register should be programmed to at least 512/32=16='10000' to conform to the DDR3 specification.

2.10.16 MC_CHANNEL_0_RCOMP_PARAMS MC_CHANNEL_1_RCOMP_PARAMS MC_CHANNEL_2_RCOMP_PARAMS

This register contains parameters that specify Rcomp timings.

Device: 4, 5, 6 Function: 0 Offset: 98h Access as a Dword			
Bit	Type	Reset Value	Description
16	RW	1	RCOMP_EN. Enable Rcomp. When set, the Integrated Memory Controller will do the programmed blocking of requests and send indications.
15:10	RW	2	RCOMP_CMD_DCLK. Delay from the start of an RCOMP command blocking period in which the command rcomp update is done. Program this field to 15 for all configurations.
9:4	RW	9	RCOMP_LENGTH. Number of Dclks during which all commands are blocked for an RCOMP update. Data RCOMP update is done on the last DCLK of this period. Program this field to 31 for all configurations.
3:0	RW	0	RCOMP_INTERVAL. Duration of interval between Rcomp in increments of tRefI. Register value is tRefI-1. For example a setting of 0 will produce an interval of tRefI.

2.10.17 MC_CHANNEL_0_ODT_PARAMS1 MC_CHANNEL_1_ODT_PARAMS1 MC_CHANNEL_2_ODT_PARAMS1

This register contains parameters that specify ODT timings. All values are in DCLK.

Device: 4, 5, 6 Function: 0 Offset: 9Ch Access as a Dword			
Bit	Type	Reset Value	Description
26:24	RW	0	TAOFD. ODT turn off delay.
23:20	RW	6	MCODT_DURATION. Controls the duration of MC ODT activation. BL/2 + 2.
19:16	RW	4	MCODT_DELAY. Controls the delay from Rd CAS to MC ODT activation. This value is tCAS-1.
15:12	RW	5	ODT_RD_DURATION. Controls the duration of Rd ODT activation. This value is BL/2 + 2.
11:8	RW	0	ODT_RD_DELAY. Controls the delay from Rd CAS to ODT activation. This value is tCAS-tWL.
7:4	RW	5	ODT_WR_DURATION. Controls the duration of Wr ODT activation. value is BL/2 + 2.
3:0	RW	0	ODT_WR_DELAY. Controls the delay from Wr CAS to ODT activation. This value is always 0.

2.10.18 MC_CHANNEL_0_ODT_PARAMS2 MC_CHANNEL_1_ODT_PARAMS2 MC_CHANNEL_2_ODT_PARAMS2

This register contains parameters that specify Forcing ODT on Specific ranks. This register is used in debug only and not during normal operation.

Device: 4, 5, 6 Function: 0 Offset: A0h Access as a Dword			
Bit	Type	Reset Value	Description
9	RW	0	MCODT_Writes. Drive MC ODT on reads and writes.
8	RW	0	FORCE_MCODT. Force MC ODT to always be asserted.
7	RW	0	FORCE_ODT7. Force ODT for Rank7 to always be asserted.
6	RW	0	FORCE_ODT6. Force ODT for Rank6 to always be asserted.
5	RW	0	FORCE_ODT5. Force ODT for Rank5 to always be asserted.
4	RW	0	FORCE_ODT4. Force ODT for Rank4 to always be asserted.
3	RW	0	FORCE_ODT3. Force ODT for Rank3 to always be asserted.
2	RW	0	FORCE_ODT2. Force ODT for Rank2 to always be asserted.
1	RW	0	FORCE_ODT1. Force ODT for Rank1 to always be asserted.
0	RW	0	FORCE_ODT0. Force ODT for Rank0 to always be asserted.

2.10.19 MC_CHANNEL_0_ODT_MATRIX_RANK_0_3_RD MC_CHANNEL_1_ODT_MATRIX_RANK_0_3_RD MC_CHANNEL_2_ODT_MATRIX_RANK_0_3_RD

This register contains the ODT activation matrix for RANKS 0 to 3 for Reads.

Device: 4, 5, 6 Function: 0 Offset: A4h Access as a Dword			
Bit	Type	Reset Value	Description
31:24	RW	1	ODT_RD3. Bit patterns driven out onto ODT pins when Rank3 is read.
23:16	RW	1	ODT_RD2. Bit patterns driven out onto ODT pins when Rank2 is read.
15:8	RW	4	ODT_RD1. Bit patterns driven out onto ODT pins when Rank1 is read.
7:0	RW	4	ODT_RDO. Bit patterns driven out onto ODT pins when Rank0 is read.

2.10.20 MC_CHANNEL_0_ODT_MATRIX_RANK_4_7_RD MC_CHANNEL_1_ODT_MATRIX_RANK_4_7_RD MC_CHANNEL_2_ODT_MATRIX_RANK_4_7_RD

This register contains the ODT activation matrix for RANKS 4 to 7 for Reads.

Device: 4, 5, 6 Function: 0 Offset: A8h Access as a Dword			
Bit	Type	Reset Value	Description
31:24	RW	1	ODT_RD7. Bit patterns driven out onto ODT pins when Rank7 is read.
23:16	RW	1	ODT_RD6. Bit patterns driven out onto ODT pins when Rank6 is read.
15:8	RW	4	ODT_RD5. Bit patterns driven out onto ODT pins when Rank5 is read.
7:0	RW	4	ODT_RD4. Bit patterns driven out onto ODT pins when Rank4 is read.

2.10.21 MC_CHANNEL_0_ODT_MATRIX_RANK_0_3_WR MC_CHANNEL_1_ODT_MATRIX_RANK_0_3_WR MC_CHANNEL_2_ODT_MATRIX_RANK_0_3_WR

This register contains the ODT activation matrix for RANKS 0 to 3 for Writes.

Device: 4, 5, 6 Function: 0 Offset: ACh Access as a Dword			
Bit	Type	Reset Value	Description
31:24	RW	9	ODT_WR3. Bit patterns driven out onto ODT pins when Rank3 is written.
23:16	RW	5	ODT_WR2. Bit patterns driven out onto ODT pins when Rank2 is written.
15:8	RW	6	ODT_WR1. Bit patterns driven out onto ODT pins when Rank1 is written.
7:0	RW	5	ODT_WRO. Bit patterns driven out onto ODT pins when Rank0 is written.

2.10.22 MC_CHANNEL_0_ODT_MATRIX_RANK_4_7_WR MC_CHANNEL_1_ODT_MATRIX_RANK_4_7_WR MC_CHANNEL_2_ODT_MATRIX_RANK_4_7_WR

This register contains the ODT activation matrix for RANKS 4 to 7 for Writes.

Device: 4, 5, 6 Function: 0 Offset: B0h Access as a Dword			
Bit	Type	Reset Value	Description
31:24	RW	9	ODT_WR7. Bit patterns driven out onto ODT pins when Rank7 is written.
23:16	RW	5	ODT_WR6. Bit patterns driven out onto ODT pins when Rank6 is written.
15:8	RW	6	ODT_WR5. Bit patterns driven out onto ODT pins when Rank5 is written.
7:0	RW	5	ODT_WR4. Bit patterns driven out onto ODT pins when Rank4 is written

2.10.23 MC_CHANNEL_0_WAQ_PARAMS MC_CHANNEL_1_WAQ_PARAMS MC_CHANNEL_2_WAQ_PARAMS

This register contains parameters that specify settings for the Write Address Queue.

Device: 4, 5, 6 Function: 0 Offset: B4h Access as a Dword			
Bit	Type	Reset Value	Description
29:25	RW	6	PRECASWRTHRESHOLD. Threshold above which Medium-Low Priority reads cannot PRE-CAS write requests.
24:20	RW	31	PARTWRTHRESHOLD. Threshold used to raise the priority of underfill requests in the scheduler. Set to 31 to disable.
19:15	RW	31	ISOCEXITTHRESHOLD. Write Major Mode ISOC Exit Threshold. When the number of writes in the WAQ drops below this threshold, the MC will exit write major mode in the presence of a read.
14:10	RW	31	ISOCENTRYTHRESHOLD. Write Major Mode ISOC Entry Threshold. When the number of writes in the WAQ exceeds this threshold, the MC will enter write major mode in the presence of a read.
9:5	RW	22	WMENTRYTHRESHOLD. Write Major Mode Entry Threshold. When the number of writes in the WAQ exceeds this threshold, the MC will enter write major mode.
4:0	RW	22	WMEXITTHRESHOLD. Write Major Mode Exit Threshold. When the number of writes in the WAQ drop below this threshold, the MC will exit write major mode.

2.10.24 MC_CHANNEL_0_SCHEDULER_PARAMS MC_CHANNEL_1_SCHEDULER_PARAMS MC_CHANNEL_2_SCHEDULER_PARAMS

These are the parameters used to control parameters within the scheduler.

Device: 4, 5, 6 Function: 0 Offset: B8h Access as a Dword			
Bit	Type	Reset Value	Description
12	RW	1	CS_FOR_CKE_TRANSITION. Specifies if chip select is to be asserted when CKE transitions with PowerDown entry/exit and SelfRefresh exit.
11	RW	0	FLOAT_EN. When set, the address and command lines will float to save power when commands are not being sent out. This setting may not work with RDIMMs.
10:6	RW	7	PRECASRDTHRESHOLD. Threshold above which Medium-Low Priority reads can PRE-CAS write requests.
5	RW	0	DISABLE_ISOC_RBC_RESERVE. When set this bit will prevent any RBC's from being reserved for ISOC.
3	RW	0	ENABLE2N. Enable 2n Timing.
2:0	RW	0	PRIORITYCOUNTER. Upper 3 MSB of 8 bit priority time out counter.

2.10.25 MC_CHANNEL_0_MAINTENANCE_OPS MC_CHANNEL_1_MAINTENANCE_OPS MC_CHANNEL_2_MAINTENANCE_OPS

This register enables various maintenance operations such as Refreshes, ZQ, RCOMP, etc.

Device: 4, 5, 6 Function: 0 Offset: BCh Access as a Dword			
Bit	Type	Reset Value	Description
12:0	RW	0	MAINT_CNTR. Value to be loaded in the maintenance counter. This counter sequences the rate to Refreshes, ZQ, RCOMP.

2.10.26 MC_CHANNEL_0_TX_BG_SETTINGS MC_CHANNEL_1_TX_BG_SETTINGS MC_CHANNEL_2_TX_BG_SETTINGS

These are the parameters used to set the Start Scheduler for TX clock crossing. This is used to send commands to the DIMMs.

The NATIVE RATIO is UCLK multiplier of BCLK = U

ALIEN RATION is DCLK multiplier of BCLK = D

PIPE DEPTH = 8 UCLK (design dependent variable)

MIN SEP DELAY = 670ps (design dependent variable, Internally this is logic delay of FIFO + clock skew between U and D)

TOTAL EFFECTIVE DELAY = PIPE DEPTH * UCLK PERIOD in ps + MIN SEP DELAY

DELAY FRACTION = (TOTAL EFFECTIVE DELAY * D) / (UCLK PERIOD in ps * G.C.D(U,D))

Determine OFFSET MULTIPLE using the equation

FLOOR ((OFFSET MULTIPLE +1) / G.C.D (U,D)) > DELAY FRACTION

OFFSET VALUE = MOD (OFFSET MULTIPLE, U) <= Final answer for OFFSET MULTIPLE

Device: 4, 5, 6 Function: 0 Offset: C0h Access as a Dword			
Bit	Type	Reset Value	Description
23:16	RW	2	OFFSET. TX offset setting.
15:8	RW	1	ALIENRATIO. Dclk ratio to BCLK. TX Alien Ratio setting.
7:0	RW	4	NATIVERATIO. Uclk ratio to BCLK. TX Native Ratio setting.

2.10.27 MC_CHANNEL_0_RX_BGF_SETTINGS MC_CHANNEL_1_RX_BGF_SETTINGS MC_CHANNEL_2_RX_BGF_SETTINGS

These are the parameters used to set the Rx clock crossing BGF.

Device: 4, 5, 6 Function: 0 Offset: C8h Access as a Dword			
Bit	Type	Reset Value	Description
26:24	RW	2	PTRSEP. RX FIFO pointer separation settings. THIS FIELD IS NOT USED BY HARDWARE. RX Pointer separation can be modified via the round trip setting (larger value causes a larger pointer separation).
23:16	RW	0	OFFSET. RX offset setting.
15:8	RW	1	ALIENRATIO. Qclk to BCLK ratio. RX Alien Ratio setting.
7:0	RW	2	NATIVERATIO. Uclk to BCLK ratio. RX Native Ratio setting.

2.10.28 MC_CHANNEL_0_EW_BGF_SETTINGS MC_CHANNEL_1_EW_BGF_SETTINGS MC_CHANNEL_2_EW_BGF_SETTINGS

These are the parameters used to set the early warning RX clock crossing BGF.

Device: 4, 5, 6 Function: 0 Offset: CCh Access as a Dword			
Bit	Type	Reset Value	Description
15:8	RW	1	ALIENRATIO. Dclk to Bclk ratio. Early warning Alien Ratio setting.

2.10.29 MC_CHANNEL_0_EW_BGF_OFFSET_SETTINGS MC_CHANNEL_1_EW_BGF_OFFSET_SETTINGS MC_CHANNEL_2_EW_BGF_OFFSET_SETTINGS

These are the parameters to set the early warning RX clock crossing BGF.

Device: 4, 5, 6 Function: 0 Offset: D0h Access as a Dword			
Bit	Type	Reset Value	Description
15:8	RW	2	EVENOFFSET. Early warning even offset setting.
7:0	RW	0	ODDOFFSET. Early warning odd offset setting.

2.10.30 MC_CHANNEL_0_ROUND_TRIP_LATENCY MC_CHANNEL_1_ROUND_TRIP_LATENCY MC_CHANNEL_2_ROUND_TRIP_LATENCY

These are the parameters to set the early warning RX clock crossing the Bubble Generator FIFO (BGF) used to go between different clocking domains. These settings provide the gearing necessary to make that clock crossing.

Device: 4, 5, 6 Function: 0 Offset: D4h Access as a Dword			
Bit	Type	Reset Value	Description
7:0	RW	0	ROUND_TRIP_LATENCY. Round trip latency for reads. Units are in UCLK. This register must be programmed with the appropriate time for read data to be returned from the pads after a READ CAS is sent to the DIMMs.



2.10.31 MC_CHANNEL_0_PAGETABLE_PARAMS1 MC_CHANNEL_1_PAGETABLE_PARAMS1 MC_CHANNEL_2_PAGETABLE_PARAMS1

These are the parameters used to control parameters for page closing policies.

Device: 4, 5, 6 Function: 0 Offset: D8h Access as a Dword			
Bit	Type	Reset Value	Description
15:8	RW	0	REQUESTCOUNTER. This field is the upper 8 MSBs of a 12-bit counter. This counter determines the window over which the page close policy is evaluated.
7:0	RW	0	ADAPTIVETIMEOUTCOUNTER. This field is the upper 8 MSBs of a 12-bit counter. This counter adapts the interval between assertions of the page close flag. For a less aggressive page close, the length of the count interval is increased and vice versa for a more aggressive page close policy.

2.10.32 MC_CHANNEL_0_PAGETABLE_PARAMS2 MC_CHANNEL_1_PAGETABLE_PARAMS2 MC_CHANNEL_2_PAGETABLE_PARAMS2

These are the parameters used to control parameters for page closing policies.

Device: 4, 5, 6 Function: 0 Offset: DCh Access as a Dword			
Bit	Type	Reset Value	Description
27	RW	0	ENABLEADAPTIVEPAGECLOSE. 1 = Enables Adaptive Page Closing.
26:18	RW	0	MINPAGECLOSELIMIT. This field is the upper 9 MSBs of a 13-bit threshold limit. When the mistake counter falls below this threshold, a less aggressive page close interval (larger) is selected.
17:9	RW	0	MAXPAGECLOSELIMIT. This field is the upper 9 bits of a 13-bit threshold limit. When the mistake counter exceeds this threshold, a more aggressive page close interval (smaller) is selected.
8:0	RW	0	MISTAKECOUNTER. This field is the upper 8 MSBs of a 12-bit counter. This counter adapts the interval between assertions of the page close flag. For a less aggressive page close, the length of the count interval is increased and vice versa for a more aggressive page close policy.

**2.10.33 MC_TX_BG_CMD_DATA_RATIO_SETTINGS_CH0
 MC_TX_BG_CMD_DATA_RATIO_SETTINGS_CH1
 MC_TX_BG_CMD_DATA_RATIO_SETTINGS_CH2**

Channel Bubble Generator ratios for CMD and DATA.

Device: 4, 5, 6 Function: 0 Offset: E0h Access as a Dword			
Bit	Type	Reset Value	Description
15:8	RW	1	ALIENRATIO. DCLK to BCLK ratio.
7:0	RW	4	NATIVERATIO. UCLK to BCLK ratio.

**2.10.34 MC_TX_BG_CMD_OFFSET_SETTINGS_CH0
 MC_TX_BG_CMD_OFFSET_SETTINGS_CH1
 MC_TX_BG_CMD_OFFSET_SETTINGS_CH2**

Integrated Memory Controller Channel Bubble Generator Offsets for CMD FIFO. The Data command FIFOs share the settings for channel 0 across all three channels. The register in Channel 0 must be programmed for all configurations.

Device: 4, 5, 6 Function: 0 Offset: E4h Access as a Dword			
Bit	Type	Reset Value	Description
9:8	RW	0	PTROFFSET. FIFO pointer offset.
7:0	RW	0	BGOFFSET. BG offset.

**2.10.35 MC_TX_BG_DATA_OFFSET_SETTINGS_CH0
 MC_TX_BG_DATA_OFFSET_SETTINGS_CH1
 MC_TX_BG_DATA_OFFSET_SETTINGS_CH2**

Integrated Memory Controller Channel Bubble Generator Offsets for DATA FIFO.

Device: 4, 5, 6 Function: 0 Offset: E8h Access as a Dword			
Bit	Type	Reset Value	Description
16:14	RW	0	RDPTROFFSET. Read FIFO pointer offset.
13:10	RW	0	WRTPTROFFSET. Write FIFO pointer offset.
9:8	RW	0	PTROFFSET. FIFO pointer offset.
7:0	RW	0	BGOFFSET. BG offset.

2.10.36 MC_CHANNEL_0_ADDR_MATCH MC_CHANNEL_1_ADDR_MATCH MC_CHANNEL_2_ADDR_MATCH

This register specifies the intended address or address range where ECC errors will be injected. It can be set to match memory address on a per channel basis. The address fields can be masked in the Mask bits. Any mask bits set to 1 will always match. To match all addresses, all of the mask bits can be set to 1. The MC_CHANNEL_X_ECC_ERROR_INJECT register can be used to set the trigger for the error injection.

Device: 4, 5, 6 Function: 0 Offset: F0h Access as a Qword			
Bit	Type	Reset Value	Description
41	RW	0	MASK_DIMM. 1 = If set, ignore DIMM address during address comparison.
40	RW	0	MASK_RANK. 1 = If set, ignore RANK address during address comparison.
39	RW	0	MASK_BANK. 1 = If set, ignore BANK address during address comparison.
38	RW	0	MASK_PAGE. If set, ignore PAGE address during address comparison.
37	RW	0	MASK_COL. 1 = If set ignore, COLUMN address during address comparison.
36	RW	0	DIMM. DIMM address for 1 or 2DPC. For 3DPC, bits 36 and 35 represent the DIMM address and bit 34 represent the RANK address.
35:34	RW	0	RANK. Rank address for 1 or 2DPC. For 3DPC, bits 36 and 35 represent the DIMM address and bit 34 represent the RANK address.
33:30	RW	0	BANK. Bank address.
29:14	RW	0	PAGE. Page address.
13:0	RW	0	COLUMN. Column address.

2.10.37 MC_CHANNEL_0_ECC_ERROR_MASK MC_CHANNEL_1_ECC_ERROR_MASK MC_CHANNEL_2_ECC_ERROR_MASK

This register contains mask bits for the memory controller and specifies at which ECC bit(s) the error injection should occur. Any bits set to a 1 will flip the corresponding ECC bit. Correctable errors can be injected by flipping 1 bit or the bits within a symbol pair (2 consecutive aligned 8-bit pairs - i.e. 7:0 and 15:8 or 23:16 and 31:24). Flipping bits in two symbol pairs will cause an uncorrectable error to be injected.

Device: 4, 5, 6 Function: 0 Offset: F8h Access as a Dword			
Bit	Type	Reset Value	Description
31:0	RW	0	ECCMASK. This field contains the 32 bits of MC ECC mask bit for half cacheline.

2.10.38 MC_CHANNEL_0_ECC_ERROR_INJECT MC_CHANNEL_1_ECC_ERROR_INJECT MC_CHANNEL_2_ECC_ERROR_INJECT

This register contains the control bits for the actual ECC error injection. This register needs to be written after writing into MC_CHANNEL_X_ECC_ERROR_MASK. The INJECT_ECC bit must be set to enable error injection. Otherwise, no error injection will take place even if the criteria programmed in the MC_CHANNEL_X_ADDR_MATCH register is met.

Device: 4, 5, 6 Function: 0 Offset: FCh Access as a Dword			
Bit	Type	Reset Value	Description
4	RW	0	INJECT_ADDR_PARITY. 1 = Forces Address Parity error injection. Bit will reset after the first injection unless REPEAT_EN is set.
3	RW	0	INJECT_ECC. 1 = Forces ECC error injection. Bit will reset after the first injection unless REPEAT_EN is set.
2:1	RW	0	MASK_HALF_CACHELINE. 11 = Inject the ECC code word for full cacheline. 10 = Inject the ECC code word for upper 32B half cacheline. 01 = Inject the ECC code word for lower 32B half cacheline. 00 = No masking will be applied.
0	RW	0	REPEAT_EN. 1 = ECC errors will be injected on the channel until the bit is cleared.



2.10.39 Error Injection Implementation

The usage model is to program the MC_CHANNEL_X_ADDR_MATCH and MC_CHANNEL_X_ECC_ERROR_MASK registers before writing the command in MC_CHANNEL_X_ECC_ERROR_INJECT register. When writing the MC_CHANNEL_X_ECC_ERROR_INJECT register, the REPEAT_EN and MASK_HALF_CACHELINE bits need to be set to the desired values.

To turn off the feature, write 0 to the MC_CHANNEL_X_ECC_ERROR_INJECT register.

Address parity error injection and ECC error injection can be done either at the same time or independently. They will both use the same MATCH settings if both are enabled.

Note: Along with the INJECT_ECC bit set, software must generate the memory traffic that matches the address location programmed in the MC_CHANNEL_X_ADDR_MATCH register as described above in order for an error injection to take place. Unless the REPEAT_EN bit is set in the MC_CHANNEL_X_ECC_ERROR_INJECT register, the memory controller will only inject the error to the first location that matches the criteria programmed in the MC_CHANNEL_X_ADDR_MATCH register.

Errors are injected on writes only. Reads will be required to detect the errors in the MC_COR_ECC_CNT_X registers. Additionally, all writes used to inject errors must be committed to memory to ensure the error is detected on subsequent reads.

2.11 Integrated Memory Controller Channel Address Registers

2.11.1 MC_DOD_CHO_0, MC_DOD_CHO_1, MC_DOD_CHO_2

Channel 0 DIMM Organization Descriptor Register.

Device: 4 Function: 1 Offset: 48h, 4Ch, 50h Access as a Dword			
Bit	Type	Reset Value	Description
12:10	RW	0	RANKOFFSET. Rank Offset for calculating RANK. This corresponds to the first logical rank on the DIMM. The rank offset is always programmed to 0 for the DIMM 0 DOD registers. (DIMM 0 rank offset is always 0.) DIMM 1 DOD rank offset is either 4 for two DIMMs per channel or 2 if there are three DIMMs per channel. DIMM2 DOD rank offset is always 4 as it is only used in three DIMMs per channel case.
9	RW	0	DIMMPRESENT. DIMM slot is populated.
8:7	RW	0	NUMBANK. Defines the number of (real, not shadow) banks on these DIMMs. 00 = Four-banked 01 = Eight-banked 10 = Sixteen-banked
6:5	RW	0	NUMRANK. Number of Ranks. Defines the number of ranks on these DIMMs. 00 = Single Ranked 01 = Double Ranked 10 = Quad Ranked
4:2	RW	0	NUMROW. Number of Rows. Defines the number of rows within these DIMMs. 000 = 2^{12} Rows 001 = 2^{13} Rows 010 = 2^{14} Rows 011 = 2^{15} Rows 100 = 2^{16} Rows
1:0	RW	0	NUMCOL. Number of Columns. Defines the number of columns within on these DIMMs. 00 = 2^{10} columns 01 = 2^{11} columns 10 = 2^{12} columns 11 = RSVD.

2.11.2 MC_DOD_CH1_0, MC_DOD_CH1_1, MC_DOD_CH1_2

Channel 1 DIMM Organization Descriptor Register.

Device: 5 Function: 1 Offset: 48h, 4Ch, 50h Access as a Dword			
Bit	Type	Reset Value	Description
12:10	RW	0	RANKOFFSET. Rank Offset for calculating RANK. This field corresponds to the first logical rank on the DIMM. The rank offset is always programmed to 0 for the DIMM 0 DOD registers. (DIMM 0 rank offset is always 0.) DIMM 1 DOD rank offset is either 4 for two DIMMs per channel or 2 if there are three DIMMs per channel. DIMM2 DOD rank offset is always 4 as it is only used in three DIMMs per channel case.
9	RW	0	DIMMPRESENT. DIMM slot is populated.
8:7	RW	0	NUMBANK. Defines the number of (real, not shadow) banks on these DIMMs. 00 = Four-banked 01 = Eight-banked 10 = Sixteen-banked
6:5	RW	0	NUMRANK. Number of Ranks. Defines the number of ranks on these DIMMs. 00 = Single Ranked 01 = Double Ranked 10 = Quad Ranked
4:2	RW	0	NUMROW. Number of Rows. Defines the number of rows within these DIMMs. 000 = 2^{12} Rows 001 = 2^{13} Rows 010 = 2^{14} Rows 011 = 2^{15} Rows 100 = 2^{16} Rows
1:0	RW	0	NUMCOL. Number of Columns. Defines the number of columns within on these DIMMs. 00 = 2^{10} columns 01 = 2^{11} columns 10 = 2^{12} columns 11 = RSVD.

2.11.3 MC_DOD_CH2_0, MC_DOD_CH2_1, MC_DOD_CH2_2

Channel 2 DIMM Organization Descriptor Register.

Device: 6 Function: 1 Offset: 48h, 4Ch, 50h Access as a Dword			
Bit	Type	Reset Value	Description
12:10	RW	0	RANKOFFSET. Rank Offset for calculating RANK. This field corresponds to the first logical rank on the DIMM. The rank offset is always programmed to 0 for the DIMM 0 DOD registers. (DIMM 0 rank offset is always 0.) DIMM 1 DOD rank offset is either 4 for two DIMMs per channel or 2 if there are three DIMMs per channel. DIMM2 DOD rank offset is always 4 as it is only used in three DIMMs per channel case.
9	RW	0	DIMMPRESENT. DIMM slot is populated.
8:7	RW	0	NUMBANK. Defines the number of (real, not shadow) banks on these DIMMs. 00 = Four-banked 01 = Eight-banked 10 = Sixteen-banked
6:5	RW	0	NUMRANK. Defines the number of ranks on these DIMMs. 00 = Single Ranked 01 = Double Ranked 10 = Quad Ranked
4:2	RW	0	NUMROW. Defines the number of rows within these DIMMs. 000 = 2^{12} Rows 001 = 2^{13} Rows 010 = 2^{14} Rows 011 = 2^{15} Rows 100 = 2^{16} Rows
1:0	RW	0	NUMCOL. Defines the number of columns within on these DIMMs. 00 = 2^{10} columns 01 = 2^{11} columns 10 = 2^{12} columns 11 = RSVD

2.11.4 MC_SAG_CHO_0, MC_SAG_CHO_1, MC_SAG_CHO_2 MC_SAG_CHO_3, MC_SAG_CHO_4, MC_SAG_CHO_5 MC_SAG_CHO_6, MC_SAG_CHO_7, MC_SAG_CH1_0 MC_SAG_CH1_1, MC_SAG_CH1_2, MC_SAG_CH1_3 MC_SAG_CH1_4, MC_SAG_CH1_5, MC_SAG_CH1_6 MC_SAG_CH1_7, MC_SAG_CH2_0, MC_SAG_CH2_1 MC_SAG_CH2_2, MC_SAG_CH2_3, MC_SAG_CH2_4 MC_SAG_CH2_5, MC_SAG_CH2_6, MC_SAG_CH2_7

Channel Segment Address Registers. For each of the 8 interleave ranges, they specify the offset between the System Address and the Memory Address and the System Address bits used for level 1 interleave, which should not be translated to Memory Address bits. Memory Address is calculated from System Address and the contents of these registers by the following algorithm:

```
m[39:16] = SystemAddress[39:16] + (sign extend {Offset[23:0]});  

m[15:6] = SystemAddress[15:6];  

If (Removed[2]) {bit 8 removed};  

If (Removed[1]) {bit 7 removed};  

If (Removed[0]) {bit 6 removed};  

MemoryAddress[36:6] = m[36:6];
```

The following table summarizes the combinations of removed bits and divide-by-3 operations for the various supported interleave configurations. All other combinations are not supported.

Note: If any of bits [8:6] are removed, the higher order bits are shifted down.

Removed [8:6]	Divide-By-3	Interleave
000	0	None
001	0	2-Way
011	0	4-Way
000	1	3-Way
001	1	6-Way

Device: 4 Function: 1 Offset: 80h, 84h, 88h, 8Ch, 90h, 94h, 98h, 9Ch Access as a Dword			
Bit	Type	Reset Value	Description
27	RW	0	DIVBY3. This bit indicates the rule is a 3 or 6 way interleave.
26:24	RW	0	REMOVED. These are the bits to be removed after offset subtraction. These bits correspond to System Address [8,7,6].
23:0	RW	0	OFFSET. This value should be subtracted from the current system address to create a contiguous address space within a channel. BITS 9:0 ARE RESERVED AND MUST ALWAYS BE SET TO 0.

2.12 Integrated Memory Controller Channel Rank Registers

2.12.1 MC_RIR_LIMIT_CHO_0, MC_RIR_LIMIT_CHO_1
MC_RIR_LIMIT_CHO_2, MC_RIR_LIMIT_CHO_3
MC_RIR_LIMIT_CHO_4, MC_RIR_LIMIT_CHO_5
MC_RIR_LIMIT_CHO_6, MC_RIR_LIMIT_CHO_7
MC_RIR_LIMIT_CH1_0, MC_RIR_LIMIT_CH1_1
MC_RIR_LIMIT_CH1_2, MC_RIR_LIMIT_CH1_3
MC_RIR_LIMIT_CH1_4, MC_RIR_LIMIT_CH1_5
MC_RIR_LIMIT_CH1_6, MC_RIR_LIMIT_CH1_7
MC_RIR_LIMIT_CH2_0, MC_RIR_LIMIT_CH2_1
MC_RIR_LIMIT_CH2_2, MC_RIR_LIMIT_CH2_3
MC_RIR_LIMIT_CH2_4, MC_RIR_LIMIT_CH2_5
MC_RIR_LIMIT_CH2_6, MC_RIR_LIMIT_CH2_7

Channel Rank Limit Range Registers.

Device: 4 Function: 2 Offset: 40h, 44h, 48h, 4Ch, 50h, 54h, 58h, 5Ch Access as a Dword			
Bit	Type	Reset Value	Description
9:0	RW	0	LIMIT. This field specifies the top of the range being mapped to the ranks specified in the MC_RIR_WAY_CH registers. The most significant bits of the lowest address in this range is one greater than the limit field in the RIR register with the next lower index. This field is compared against MA[37:28].



2.12.2 MC_RIR_WAY_CHO_0, MC_RIR_WAY_CHO_1 MC_RIR_WAY_CHO_2, MC_RIR_WAY_CHO_3 MC_RIR_WAY_CHO_4, MC_RIR_WAY_CHO_5 MC_RIR_WAY_CHO_6, MC_RIR_WAY_CHO_7 MC_RIR_WAY_CHO_8, MC_RIR_WAY_CHO_9 MC_RIR_WAY_CHO_10, MC_RIR_WAY_CHO_11 MC_RIR_WAY_CHO_12, MC_RIR_WAY_CHO_13 MC_RIR_WAY_CHO_14, MC_RIR_WAY_CHO_15 MC_RIR_WAY_CHO_16, MC_RIR_WAY_CHO_17 MC_RIR_WAY_CHO_18, MC_RIR_WAY_CHO_19 MC_RIR_WAY_CHO_20, MC_RIR_WAY_CHO_21 MC_RIR_WAY_CHO_22, MC_RIR_WAY_CHO_23 MC_RIR_WAY_CHO_24, MC_RIR_WAY_CHO_25 MC_RIR_WAY_CHO_26, MC_RIR_WAY_CHO_27 MC_RIR_WAY_CHO_28, MC_RIR_WAY_CHO_29 MC_RIR_WAY_CHO_30, MC_RIR_WAY_CHO_31

Channel Rank Interleave Way Range Registers. These registers allow the user to define the ranks and offsets that apply to the ranges defined by the LIMIT in the MC_RIR_LIMIT_CH registers. The mappings are as follows:

RIR_LIMIT_CH{chan}[0] -> RIR_WAY_CH{chan}[3:0]
 RIR_LIMIT_CH{chan}[1] -> RIR_WAY_CH{chan}[7:6]
 RIR_LIMIT_CH{chan}[2] -> RIR_WAY_CH{chan}[11:10]
 RIR_LIMIT_CH{chan}[3] -> RIR_WAY_CH{chan}[15:14]
 RIR_LIMIT_CH{chan}[4] -> RIR_WAY_CH{chan}[19:18]
 RIR_LIMIT_CH{chan}[5] -> RIR_WAY_CH{chan}[23:22]
 RIR_LIMIT_CH{chan}[6] -> RIR_WAY_CH{chan}[27:26]
 RIR_LIMIT_CH{chan}[7] -> RIR_WAY_CH{chan}[31:28]

Device: 4 Function: 2 Offset: 80h, 84h, 88h, 8Ch, 90h, 94h, 98h, 9Ch, A0h, A4h, A8h, ACh, B0h, B4h, B8h, BCh, C0h, C4h, C8h, CCh, D0h, D4h, D8h, DCh, E0h, E4h, E8h, ECh, F0h, F4h, F8h, FCh Access as a Dword			
Bit	Type	Reset Value	Description
13:4	RW	0	OFFSET. This field defines the offset used in the rank interleave. This is a 2's complement value.
3:0	RW	0	RANK. This field defines which rank participates in WAY(n). If MC_CONTROL.CLOSED_PAGE=1, this field defines the DRAM rank selected when MemoryAddress[7:6]=(n). If MMC_CONTROL.CLOSED_PAGE=0, this field defines which rank is selected when MemoryAddress[13:12]=(n). (n) is the instantiation of the register. This field is organized by physical rank. Bits [3:2] are the encoded DIMM ID(slot). Bits [1:0] are the rank within that DIMM.

**2.12.3 MC_RIR_WAY_CH1_0, MC_RIR_WAY_CH1_1
 MC_RIR_WAY_CH1_2, MC_RIR_WAY_CH1_3
 MC_RIR_WAY_CH1_4, MC_RIR_WAY_CH1_5
 MC_RIR_WAY_CH1_6, MC_RIR_WAY_CH1_7
 MC_RIR_WAY_CH1_8, MC_RIR_WAY_CH1_9
 MC_RIR_WAY_CH1_10, MC_RIR_WAY_CH1_11
 MC_RIR_WAY_CH1_12, MC_RIR_WAY_CH1_13
 MC_RIR_WAY_CH1_14, MC_RIR_WAY_CH1_15
 MC_RIR_WAY_CH1_16, MC_RIR_WAY_CH1_17
 MC_RIR_WAY_CH1_18, MC_RIR_WAY_CH1_19
 MC_RIR_WAY_CH1_20, MC_RIR_WAY_CH1_21
 MC_RIR_WAY_CH1_22, MC_RIR_WAY_CH1_23
 MC_RIR_WAY_CH1_24, MC_RIR_WAY_CH1_25
 MC_RIR_WAY_CH1_26, MC_RIR_WAY_CH1_27
 MC_RIR_WAY_CH1_28, MC_RIR_WAY_CH1_29
 MC_RIR_WAY_CH1_30, MC_RIR_WAY_CH1_31**

Channel Rank Interleave Way Range Registers. These registers allow the user to define the ranks and offsets that apply to the ranges defined by the LIMIT in the MC_RIR_LIMIT_CH registers. The mappings are as follows:

RIR_LIMIT_CH{chan}[0] -> RIR_WAY_CH{chan}[3:0]
 RIR_LIMIT_CH{chan}[1] -> RIR_WAY_CH{chan}[7:6]
 RIR_LIMIT_CH{chan}[2] -> RIR_WAY_CH{chan}[11:10]
 RIR_LIMIT_CH{chan}[3] -> RIR_WAY_CH{chan}[15:14]
 RIR_LIMIT_CH{chan}[4] -> RIR_WAY_CH{chan}[19:18]
 RIR_LIMIT_CH{chan}[5] -> RIR_WAY_CH{chan}[23:22]
 RIR_LIMIT_CH{chan}[6] -> RIR_WAY_CH{chan}[27:26]
 RIR_LIMIT_CH{chan}[7] -> RIR_WAY_CH{chan}[31:28]

Device: 5 Function: 2 Offset: 80h, 84h, 88h, 8Ch, 90h, 94h, 98h, 9Ch, A0h, A4h, A8h, ACh, B0h, B4h, B8h, BCh, C0h, C4h, C8h, CCh, D0h, D4h, D8h, DCh, E0h, E4h, E8h, ECh, F0h, F4h, F8h, FCh Access as a Dword			
Bit	Type	Reset Value	Description
13:4	RW	0	OFFSET. This field defines the offset used in the rank interleave. This is a 2's complement value.
3:0	RW	0	RANK. This field defines which rank participates in WAY(n). If MC_CONTROL.CLOSED_PAGE=1, this field defines the DRAM rank selected when MemoryAddress[7:6]=(n). If MC_CONTROL.CLOSED_PAGE=0, this field defines which rank is selected when MemoryAddress[13:12]=(n). (n) is the instantiation of the register. This field is organized by physical rank. Bits [3:2] are the encoded DIMM ID(slot). Bits [1:0] are the rank within that DIMM.



2.12.4 MC_RIR_WAY_CH2_0, MC_RIR_WAY_CH2_1 MC_RIR_WAY_CH2_2, MC_RIR_WAY_CH2_3 MC_RIR_WAY_CH2_4, MC_RIR_WAY_CH2_5 MC_RIR_WAY_CH2_6, MC_RIR_WAY_CH2_7 MC_RIR_WAY_CH2_8, MC_RIR_WAY_CH2_9 MC_RIR_WAY_CH2_10, MC_RIR_WAY_CH2_11 MC_RIR_WAY_CH2_12, MC_RIR_WAY_CH2_13 MC_RIR_WAY_CH2_14, MC_RIR_WAY_CH2_15 MC_RIR_WAY_CH2_16, MC_RIR_WAY_CH2_17 MC_RIR_WAY_CH2_18, MC_RIR_WAY_CH2_19 MC_RIR_WAY_CH2_20, MC_RIR_WAY_CH2_21 MC_RIR_WAY_CH2_22, MC_RIR_WAY_CH2_23 MC_RIR_WAY_CH2_24, MC_RIR_WAY_CH2_25 MC_RIR_WAY_CH2_26, MC_RIR_WAY_CH2_27 MC_RIR_WAY_CH2_28, MC_RIR_WAY_CH2_29 MC_RIR_WAY_CH2_30, MC_RIR_WAY_CH2_31

Channel Rank Interleave Way Range Registers. These registers allow the user to define the ranks and offsets that apply to the ranges defined by the LIMIT in the MC_RIR_LIMIT_CH registers. The mappings are as follows:

RIR_LIMIT_CH{chan}[0] -> RIR_WAY_CH{chan}[3:0]
 RIR_LIMIT_CH{chan}[1] -> RIR_WAY_CH{chan}[7:6]
 RIR_LIMIT_CH{chan}[2] -> RIR_WAY_CH{chan}[11:10]
 RIR_LIMIT_CH{chan}[3] -> RIR_WAY_CH{chan}[15:14]
 RIR_LIMIT_CH{chan}[4] -> RIR_WAY_CH{chan}[19:18]
 RIR_LIMIT_CH{chan}[5] -> RIR_WAY_CH{chan}[23:22]
 RIR_LIMIT_CH{chan}[6] -> RIR_WAY_CH{chan}[27:26]
 RIR_LIMIT_CH{chan}[7] -> RIR_WAY_CH{chan}[31:28]

Device: 6 Function: 2 Offset: 80h, 84h, 88h, 8Ch, 90h, 94h, 98h, 9Ch, A0h, A4h, A8h, ACh, B0h, B4h, B8h, BCh, C0h, C4h, C8h, CCh, D0h, D4h, D8h, DCh, E0h, E4h, E8h, ECh, F0h, F4h, F8h, FCh Access as a Dword			
Bit	Type	Reset Value	Description
13:4	RW	0	OFFSET. This field defines the offset used in the rank interleave. This is a 2's complement value.
3:0	RW	0	RANK. This field defines which rank participates in WAY(n). If MC_CONTROL.CLOSED_PAGE=1, this field defines the DRAM rank selected when MemoryAddress[7:6]=(n). If MC_CONTROL.CLOSED_PAGE=0, this field defines which rank is selected when MemoryAddress[13:12]=(n). (n) is the instantiation of the register. This field is organized by physical rank. Bits [3:2] are the encoded DIMM ID(slot). Bits [1:0] are the rank within that DIMM.

2.13 Memory Thermal Control

2.13.1 MC_THERMAL_CONTROLO MC_THERMAL_CONTROL1 MC_THERMAL_CONTROL2

Controls for the Integrated Memory Controller thermal throttle logic for each channel.

Device: 4, 5, 6 Function: 3 Offset: 48h Access as a Dword			
Bit	Type	Reset Value	Description
2	RW	1	APPLY_SAFE. Enable the application of safe values while MC_THERMAL_PARAMS_B.SAFE_INTERVAL is exceeded.
1:0	RW	0	THROTTLE_MODE. S elects throttling mode. 00 = Throttle disabled 01 = Open Loop: Throttle when Virtual Temperature is greater than MC_THROTTLE_OFFSET. 10 = Closed Loop: Throttle when MC_CLOSED_LOOP.THROTTLE_NOW is set. 11 = Closed Loop: Throttle when MC_DDR_THERM_COMMAND.THROTTLE is set and the MC_DDR_THERM pin is asserted OR OLTT will be implemented (Condition 1).

2.13.2 MC_THERMAL_STATUS0 MC_THERMAL_STATUS1 MC_THERMAL_STATUS2

Status registers for the thermal throttling logic for each channel.

Device: 4, 5, 6 Function: 3 Offset: 4Ch Access as a Dword			
Bit	Type	Reset Value	Description
29:4	RO	0	CYCLES_THROTTLED. The number of throttle cycles, in increments of 256 Dclks, triggered in any rank in the last SAFE_INTERVAL number of ZQs.
3:0	RO	0	RANK_TEMP. The bit specifies whether the rank is above throttling threshold.

2.13.3 MC_THERMAL_DEFEATURE0 MC_THERMAL_DEFEATURE1 MC_THERMAL_DEFEATURE2

Thermal Throttle defeature register for each channel.

Device: 4, 5, 6 Function: 3 Offset: 50h Access as a Dword			
Bit	Type	Reset Value	Description
0	RW1S	0	THERM_REG_LOCK. When set, no further modification of all thermal throttle registers are allowed. This bit must be set to the same value for all channels.

2.13.4 MC_THERMAL_PARAMS_A0 MC_THERMAL_PARAMS_A1 MC_THERMAL_PARAMS_A2

Parameters used by Open Loop Throughput Throttling (OLTT) and Closed Loop Thermal Throttling (CLTT).

Device: 4, 5, 6 Function: 3 Offset: 60h Access as a Dword			
Bit	Type	Reset Value	Description
31:24	RW	0	CKE_ASSERT_ENERGY. Energy of having CKE asserted when no command is issued.
23:16	RW	0	CKE_DEASSERT_ENERGY. Energy of having CKE de-asserted when no command is issued.
15:8	RW	0	WRCMD_ENERGY. Energy of a write including data transfer.
7:0	RW	0	RDCMD_ENERGY. Energy of a read including data transfer.

2.13.5 MC_THERMAL_PARAMS_B0 MC_THERMAL_PARAMS_B1 MC_THERMAL_PARAMS_B2

Parameters used by the thermal throttling logic.

Device: 4, 5, 6 Function: 3 Offset: 64h Access as a Dword			
Bit	Type	Reset Value	Description
31:26	RW	1	SAFE_INTERVAL. Safe values for cooling coefficient and duty cycle will be applied while the SAFE_INTERVAL is exceeded. This interval is the number of ZQ intervals since the last time the MC_COOLING_COEF or MC_CLOSED_LOOP registers have been written. A register to write to MC_COOLING_COEF or MC_CLOSED_LOOP will re-apply the normal MC_COOLING_COEF and MC_CLOSED_LOOP.MIN_THROTTLE_DUTY_CYC values. The register value written need not be different; writing the current value will suffice. The MC_THERMAL_STATUS.CYCLES_THROTTLED field is reloaded when the number of ZQ intervals exceeds this value. This field must not be programmed to 0; this value is illegal.
25:16	RW	255	SAFE_DUTY_CYC. This value replaces MC_CLOSED_LOOP.MIN_THROTTLE_DUTY_CYC while the MC_THERMAL_PARAMS_B.SAFE_INTERVAL is exceeded.
15:8	RW	1	SAFE_COOL_COEF. This value replaces MC_COOLING_COEF while the THERMAL_PARAMS_B.SAFE_INTERVAL is exceeded.
7:0	RW	0	ACTCMD_ENERGY. Energy of an Activate/Precharge Cycle.

2.13.6 MC_COOLING_COEFO MC_COOLING_COEF1 MC_COOLING_COEF2

Heat removed from DRAM 8 DCLKs. This should be scaled relative to the per command weights and the initial value of the throttling threshold. This includes idle command and refresh energies. If 2X refresh is supported, the worst case of 2X refresh must be assumed.

When there are more than 4 ranks attached to the channel, the thermal throttle logic is shared.

Device: 4, 5, 6 Function: 3 Offset: 80h Access as a Dword			
Bit	Type	Reset Value	Description
31:24	RW	255	RANK3. Rank 3 Cooling Coefficient.
23:16	RW	255	RANK2. Rank 2 Cooling Coefficient.
15:8	RW	255	RANK1. Rank 1 Cooling Coefficient.
7:0	RW	255	RANK0. Rank 0 Cooling Coefficient.

2.13.7 MC_CLOSED_LOOP0 MC_CLOSED_LOOP1 MC_CLOSED_LOOP2

This register controls the closed loop thermal response of the DRAM thermal throttle logic. It supports immediate thermal throttle and 2X refresh. In addition, the register is used to configure the throttling duty cycle.

Device: 4, 5, 6 Function: 3 Offset: 84h Access as a Dword			
Bit	Type	Reset Value	Description
17:8	RW	64	MIN_THROTTLE_DUTY_CYC. This parameter represents the minimum number of DCLKs of operation allowed after throttling. In order to provide actual command opportunities, the number of clocks between CKE de-assertion and first command should be considered.
4	RW	0	REF_2X_NOW. Direct control of dynamic 2X refresh if MC_THERMAL_CONTROL.THROTTLE_MODE = 2.
3:0	RW	0	THROTTLE_NOW. Throttler Vector to directly control throttling if MC_THERMAL_CONTROL.THROTTLE_MODE = 2.

2.13.8 MC_THROTTLE_OFFSET0 MC_THROTTLE_OFFSET1 MC_THROTTLE_OFFSET2

Compared against bits [36:29] of virtual temperature of each rank stored in RANK_VIRTUAL_TEMP to determine the throttle point. Recommended value for each rank is 255.

When there are more than 4 ranks attached to the channel, the thermal throttle logic is shared.

Device: 4, 5, 6 Function: 3 Offset: 88h Access as a Dword			
Bit	Type	Reset Value	Description
31:24	RW	0	RANK3. Rank 3 throttle offset.
23:16	RW	0	RANK2. Rank 2 throttle offset.
15:8	RW	0	RANK1. Rank 1 throttle offset.
7:0	RW	0	RANK0. Rank 0 throttle offset.

2.13.9 MC_RANK_VIRTUAL_TEMPO MC_RANK_VIRTUAL_TEMP1 MC_RANK_VIRTUAL_TEMP2

This register contains the 8 most significant bits [37:30] of the virtual temperature of each rank. The difference between the virtual temperature and the sensor temperature can be used to determine how fast fan speed should be increased. The value stored is right shifted one bit to the right with respect to the corresponding MC_Throttle_Offset register value. For example when a rank throttle offset is set to 40h, the value read from the corresponding MC_RANK_VIRTUAL_TEMP register is 20h.

When there are more than 4 ranks attached to the channel, the thermal throttle logic is shared.

Device: 4, 5, 6 Function: 3 Offset: 98h Access as a Dword			
Bit	Type	Reset Value	Description
31:24	RO	0	RANK3. Rank 3 virtual temperature.
23:16	RO	0	RANK2. Rank 2 virtual temperature.
15:8	RO	0	RANK1. Rank 1 virtual temperature.
7:0	RO	0	RANK0. Rank 0 virtual temperature.

2.13.10 MC_DDR_THERM_COMMAND0 MC_DDR_THERM_COMMAND1 MC_DDR_THERM_COMMAND2

This register contains the command portion of the DDR_THERM# functionality as described in the processor datasheet (i.e., what an assertion of the pin does).

Device: 4, 5, 6 Function: 3 Offset: 9Ch Access as a Dword			
Bit	Type	Reset Value	Description
3	RW	0	THROTTLE. Force throttling when DDR_THERM# pin is asserted.
2	RW	0	Reserved
1	RW	0	DISABLE_EXTTS. Response to DDR_THERM# pin is disabled. ASSERTION and DEASSERTION fields in the register MC_DDR_THERM_STATUS are frozen.
0	RW	0	LOCK. When set, all bits in this register are RO and cannot be written.

2.13.11 MC_DDR_THERM_STATUS0 MC_DDR_THERM_STATUS1 MC_DDR_THERM_STATUS2

This register contains the status portion of the DDR_THERM# functionality as described in the processor datasheet (i.e., what is happening or has happened with respect to the pin).

Device: 4, 5, 6 Function: 3 Offset: A4h Access as a Dword			
Bit	Type	Reset Value	Description
2	RO	0	ASSERTION. An assertion edge was seen on DDR_THERM#. Write-1-to-clear.
1	RO	0	DEASSERTION. A de-assertion edge was seen on DDR_THERM#. Write-1-to-clear.
0	RO	0	STATE. Present logical state of DDR_THERM# bit. This is a static indication of the pin, and may be several clocks out of date due to the delay between the pin and the signal. STATE = 0 means DDR_THERM# is deasserted STATE = 1 means DDR_THERM# is asserted

2.14 Integrated Memory Controller Miscellaneous Registers

2.14.1 MC_DIMM_CLK_RATIO_STATUS

This register contains status information about DIMM clock ratio.

Device: 3 Function: 4 Offset: 50h Access as a Dword			
Bit	Type	Reset Value	Description
28:24	RO	0	MAX_RATIO. Maximum ratio allowed by the part. Value = Qclk 00000 = RSVD 00110 = 800 MHz 01000 = 1066 MHz 01010 = 1333 MHz
4:0	RO	0	QCLK_RATIO. Current ratio of Qclk. Value = Qclk. 00000 = RSVD 00110 = 800 MHz 01000 = 1066 MHz 01010 = 1333 MHz

2.14.2 MC_DIMM_CLK_RATIO

This register is for the Requested DIMM clock ratio (Qclk). This is the data rate going to the DIMM. The clock sent to the DIMM is 1/2 of QCLK rate.

Device: 3 Function: 4 Offset: 54h Access as a Dword			
Bit	Type	Reset Value	Description
4:0	RW	6	QCLK_RATIO. Requested ratio of Qclk/Bclk. 00000 = RSVD 00110 = 800 MHz 01000 = 1066 MHz 01010 = 1333 MHz

S

Mouser Electronics

Authorized Distributor

Click to View Pricing, Inventory, Delivery & Lifecycle Information:

[Intel:](#)

[BX80601W3570](#)