

Intel[®] Xeon[®] Processor 5500 Series

Datasheet, Volume 2

April 2009

Order Number: 321322-002



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 5500 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.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

I²C is a two-wire communications bus/protocol developed by Philips. SMBus is a subset of the I²C bus/protocol and was developed by Intel. Implementations of the I²C bus/protocol may require licenses from various entities, including Philips Electronics N.V. and North American Philips Corporation.

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	15
1.1	Terminology	15
1.1.1	Processor Terminology	15
1.2	References	17
2	Register Description	19
2.1	Register Terminology	19
2.2	Platform Configuration Structure	20
2.3	Device Mapping	21
2.4	Detailed Configuration Space Maps	23
2.5	PCI Standard Registers	45
2.5.1	VID - Vendor Identification Register	45
2.5.2	DID - Device Identification Register	45
2.5.3	RID - Revision Identification Register	46
2.5.4	CCR - Class Code Register	46
2.5.5	HDR - Header Type Register	47
2.5.6	SID/SVID - Subsystem Identity/Subsystem Vendor Identification Register	47
2.5.7	PCICMD - Command Register	48
2.5.8	PCISTS - PCI Status Register	49
2.6	Generic Non-core Registers	50
2.6.1	MAXREQUEST_LC	50
2.6.2	MAXREQUEST_LS	51
2.6.3	MAXREQUEST_LL	51
2.6.4	MAX_RTIDS	51
2.6.5	DESIRED_CORES	52
2.6.6	MEMLOCK_STATUS	52
2.6.7	MC_CFG_CONTROL	53
2.6.8	POWER_CNTRL_ERR_STATUS	53
2.6.9	CURRENT_UCLK_RATIO	54
2.6.10	MIRROR_PORT_CTL	55
2.6.11	MIP_PH_CTR_LO	55
	MIP_PH_CTR_L1	55
2.6.12	MIP_PH_PRT_LO	56
	MIP_PH_PRT_L1	56
2.7	SAD - System Address Decoder Registers	56
2.7.1	SAD_PAM0123	56
2.7.2	SAD_PAM456	58
2.7.3	SAD_HEN	59
2.7.4	SAD_SMRAM	59
2.7.5	SAD_PCIEXBAR	60
2.7.6	SAD_DRAM_RULE_0	60
	SAD_DRAM_RULE_1	60
	SAD_DRAM_RULE_2	60
	SAD_DRAM_RULE_3	60
	SAD_DRAM_RULE_4	60
	SAD_DRAM_RULE_5	60
	SAD_DRAM_RULE_6	60
	SAD_DRAM_RULE_7	60



2.7.7	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	61
2.8	Intel QPI Link Registers	61
2.8.1	QPI_QPILCP_L0 QPI_QPILCP_L1	61
2.8.2	QPI_QPILCL_L0 QPI_QPILCL_L1	62
2.8.3	QPI_QPILS_L0 QPI_QPILS_L1	63
2.8.4	QPI_DEF_RMT_VN_CREDITS_L0 QPI_DEF_RMT_VN_CREDITS_L1	63
2.8.5	QPI_RMT_QPILP0_STAT_L0 QPI_RMT_QPILP0_STAT_L1	63
2.8.6	QPI_RMT_QPILP1_STAT_L0 QPI_RMT_QPILP1_STAT_L1	64
2.8.7	QPI_RMT_QPILP2_STAT_L0 QPI_RMT_QPILP2_STAT_L1	64
2.8.8	QPI_RMT_QPILP3_STAT_L0 QPI_RMT_QPILP3_STAT_L1	65
2.9	Intel QPI Physical Layer Registers	66
2.9.1	QPI_0_PH_CPR QPI_1_PH_CPR	66
2.9.2	QPI_0_PH_CTR QPI_1_PH_CTR	67
2.9.3	QPI_0_PH_PIS QPI_1_PH_PIS	68
2.9.4	QPI_0_PH_PTV QPI_1_PH_PTV	69
2.9.5	QPI_0_PH_LDC QPI_1_PH_LDC	69
2.9.6	QPI_0_PH_PRT QPI_1_PH_PRT	70
2.9.7	QPI_0_PH_PMRO QPI_1_PH_PMRO	70
2.9.8	QPI_0_EP_SR QPI_1_EP_SR	71
2.9.9	QPI_0_EP_MCTR QPI_1_EP_MCTR	71
2.10	Intel QPI Miscellaneous Registers	72
2.10.1	QPI_0_PLL_STATUS QPI_1_PLL_STATUS	72
2.10.2	QPI_0_PLL_RATIO QPI_1_PLL_RATIO	72
2.11	Integrated Memory Controller Control Registers	73
2.11.1	MC_CONTROL	73
2.11.2	MC_STATUS	74
2.11.3	MC_SMI_DIMM_ERROR_STATUS	74
2.11.4	MC_SMI_CNTRL	75
2.11.5	MC_RESET_CONTROL	76
2.11.6	MC_CHANNEL_MAPPER	76
2.11.7	MC_MAX_DOD	77
2.11.8	MC_RD_CRDT_INIT	77
2.11.9	MC_CRDT_WR_THLD	78



2.11.10	MC_SCRUBADDR_LO	79
2.11.11	MC_SCRUBADDR_HI	79
2.12	TAD - Target Address Decoder Registers	80
2.12.1	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.....	80
2.12.2	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.....	81
2.13	Integrated Memory Controller RAS Registers.....	82
2.13.1	MC_SSRCONTROL	82
2.13.2	MC_SCRUB_CONTROL	83
2.13.3	MC_RAS_ENABLES.....	83
2.13.4	MC_RAS_STATUS	83
2.13.5	MC_SSRSTATUS	84
2.13.6	MC_COR_ECC_CNT_0 MC_COR_ECC_CNT_1 MC_COR_ECC_CNT_2 MC_COR_ECC_CNT_3 MC_COR_ECC_CNT_4 MC_COR_ECC_CNT_5.....	84
2.14	Integrated Memory Controller Test Registers.....	85
2.14.1	MC_TEST_ERR_RCV1	85
2.14.2	MC_TEST_ERR_RCV0	85
2.14.3	MC_TEST_PH_CTR	86
2.14.4	MC_TEST_PH_PIS	86
2.14.5	MC_TEST_PAT_GCTR	86
2.14.6	MC_TEST_PAT_BA	87
2.14.7	MC_TEST_PAT_IS	87
2.14.8	MC_TEST_PAT_DCD.....	87
2.15	Integrated Memory Controller Channel Control Registers	88
2.15.1	MC_CHANNEL_0_DIMM_RESET_CMD MC_CHANNEL_1_DIMM_RESET_CMD MC_CHANNEL_2_DIMM_RESET_CMD.....	88
2.15.2	MC_CHANNEL_0_DIMM_INIT_CMD MC_CHANNEL_1_DIMM_INIT_CMD MC_CHANNEL_2_DIMM_INIT_CMD.....	88
2.15.3	MC_CHANNEL_0_DIMM_INIT_PARAMS MC_CHANNEL_1_DIMM_INIT_PARAMS MC_CHANNEL_2_DIMM_INIT_PARAMS	89
2.15.4	MC_CHANNEL_0_DIMM_INIT_STATUS MC_CHANNEL_1_DIMM_INIT_STATUS MC_CHANNEL_2_DIMM_INIT_STATUS.....	91
2.15.5	MC_CHANNEL_0_DDR3CMD MC_CHANNEL_1_DDR3CMD MC_CHANNEL_2_DDR3CMD	92
2.15.6	MC_CHANNEL_0_REFRESH_THROTTLE_SUPPORT MC_CHANNEL_1_REFRESH_THROTTLE_SUPPORT MC_CHANNEL_2_REFRESH_THROTTLE_SUPPORT.....	93



2.15.7	MC_CHANNEL_0_MRS_VALUE_0_1 MC_CHANNEL_1_MRS_VALUE_0_1 MC_CHANNEL_2_MRS_VALUE_0_1	93
2.15.8	MC_CHANNEL_0_MRS_VALUE_2 MC_CHANNEL_1_MRS_VALUE_2 MC_CHANNEL_2_MRS_VALUE_2	94
2.15.9	MC_CHANNEL_0_RANK_PRESENT MC_CHANNEL_1_RANK_PRESENT MC_CHANNEL_2_RANK_PRESENT	94
2.15.10	MC_CHANNEL_0_RANK_TIMING_A MC_CHANNEL_1_RANK_TIMING_A MC_CHANNEL_2_RANK_TIMING_A	95
2.15.11	MC_CHANNEL_0_RANK_TIMING_B MC_CHANNEL_1_RANK_TIMING_B MC_CHANNEL_2_RANK_TIMING_B	98
2.15.12	MC_CHANNEL_0_BANK_TIMING MC_CHANNEL_1_BANK_TIMING MC_CHANNEL_2_BANK_TIMING	99
2.15.13	MC_CHANNEL_0_REFRESH_TIMING MC_CHANNEL_1_REFRESH_TIMING MC_CHANNEL_2_REFRESH_TIMING	99
2.15.14	MC_CHANNEL_0_CKE_TIMING MC_CHANNEL_1_CKE_TIMING MC_CHANNEL_2_CKE_TIMING	100
2.15.15	MC_CHANNEL_0_ZQ_TIMING MC_CHANNEL_1_ZQ_TIMING MC_CHANNEL_2_ZQ_TIMING	100
2.15.16	MC_CHANNEL_0_RCOMP_PARAMS MC_CHANNEL_1_RCOMP_PARAMS MC_CHANNEL_2_RCOMP_PARAMS	101
2.15.17	MC_CHANNEL_0_ODT_PARAMS1 MC_CHANNEL_1_ODT_PARAMS1 MC_CHANNEL_2_ODT_PARAMS1	101
2.15.18	MC_CHANNEL_0_ODT_PARAMS2 MC_CHANNEL_1_ODT_PARAMS2 MC_CHANNEL_2_ODT_PARAMS2	102
2.15.19	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	102
2.15.20	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	103
2.15.21	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	103
2.15.22	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	103
2.15.23	MC_CHANNEL_0_WAQ_PARAMS MC_CHANNEL_1_WAQ_PARAMS MC_CHANNEL_2_WAQ_PARAMS	104
2.15.24	MC_CHANNEL_0_SCHEDULER_PARAMS MC_CHANNEL_1_SCHEDULER_PARAMS MC_CHANNEL_2_SCHEDULER_PARAMS	104
2.15.25	MC_CHANNEL_0_MAINTENANCE_OPS MC_CHANNEL_1_MAINTENANCE_OPS MC_CHANNEL_2_MAINTENANCE_OPS	105
2.15.26	MC_CHANNEL_0_TX_BG_SETTINGS MC_CHANNEL_1_TX_BG_SETTINGS MC_CHANNEL_2_TX_BG_SETTINGS	105



2.15.27	MC_CHANNEL_0_RX_BGF_SETTINGS MC_CHANNEL_1_RX_BGF_SETTINGS MC_CHANNEL_2_RX_BGF_SETTINGS	106
2.15.28	MC_CHANNEL_0_EW_BGF_SETTINGS MC_CHANNEL_1_EW_BGF_SETTINGS MC_CHANNEL_2_EW_BGF_SETTINGS	106
2.15.29	MC_CHANNEL_0_EW_BGF_OFFSET_SETTINGS MC_CHANNEL_1_EW_BGF_OFFSET_SETTINGS MC_CHANNEL_2_EW_BGF_OFFSET_SETTINGS	106
2.15.30	MC_CHANNEL_0_ROUND_TRIP_LATENCY MC_CHANNEL_1_ROUND_TRIP_LATENCY MC_CHANNEL_2_ROUND_TRIP_LATENCY	107
2.15.31	MC_CHANNEL_0_PAGETABLE_PARAMS1 MC_CHANNEL_1_PAGETABLE_PARAMS1 MC_CHANNEL_2_PAGETABLE_PARAMS1	107
2.15.32	MC_TX_BG_CMD_DATA_RATIO_SETTINGS_CH0 MC_TX_BG_CMD_DATA_RATIO_SETTINGS_CH1 MC_TX_BG_CMD_DATA_RATIO_SETTINGS_CH2	107
2.15.33	MC_TX_BG_CMD_OFFSET_SETTINGS_CH0 MC_TX_BG_CMD_OFFSET_SETTINGS_CH1 MC_TX_BG_CMD_OFFSET_SETTINGS_CH2	108
2.15.34	MC_TX_BG_DATA_OFFSET_SETTINGS_CH0 MC_TX_BG_DATA_OFFSET_SETTINGS_CH1 MC_TX_BG_DATA_OFFSET_SETTINGS_CH2	108
2.16	Integrated Memory Controller Channel Address Registers	109
2.16.1	MC_DOD_CH0_0 MC_DOD_CH0_1 MC_DOD_CH0_2	109
2.16.2	MC_DOD_CH1_0 MC_DOD_CH1_1 MC_DOD_CH1_2	110
2.16.3	MC_DOD_CH2_0 MC_DOD_CH2_1 MC_DOD_CH2_2	111
2.16.4	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	112
2.17	Integrated Memory Controller Channel Rank Registers	113



2.17.1	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.....	113
2.17.2	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.....	114



2.17.3	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	115



2.17.4	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.....	117
2.18	Memory Thermal Control	118
2.18.1	MC_THERMAL_CONTROLO	
	MC_THERMAL_CONTROL1	
	MC_THERMAL_CONTROL2	118
2.18.2	MC_THERMAL_STATUS0	
	MC_THERMAL_STATUS1	
	MC_THERMAL_STATUS2.....	119
2.18.3	MC_THERMAL_DEFEATURE0	
	MC_THERMAL_DEFEATURE1	
	MC_THERMAL_DEFEATURE2	119
2.18.4	MC_THERMAL_PARAMS_A0	
	MC_THERMAL_PARAMS_A1	
	MC_THERMAL_PARAMS_A2	119
2.18.5	MC_THERMAL_PARAMS_B0	
	MC_THERMAL_PARAMS_B1	
	MC_THERMAL_PARAMS_B2	120
2.18.6	MC_COOLING_COEF0	
	MC_COOLING_COEF1	
	MC_COOLING_COEF2	120
2.18.7	MC_CLOSED_LOOP0	
	MC_CLOSED_LOOP1	
	MC_CLOSED_LOOP2	121
2.18.8	MC_THROTTLE_OFFSET0	
	MC_THROTTLE_OFFSET1	
	MC_THROTTLE_OFFSET2.....	121
2.18.9	MC_RANK_VIRTUAL_TEMPO	
	MC_RANK_VIRTUAL_TEMP1	
	MC_RANK_VIRTUAL_TEMP2	122



2.18.10	MC_DDR_THERM_COMMAND0 MC_DDR_THERM_COMMAND1 MC_DDR_THERM_COMMAND2	122
2.18.11	MC_DDR_THERM_STATUS0 MC_DDR_THERM_STATUS1 MC_DDR_THERM_STATUS2	123
2.19	Integrated Memory Controller Miscellaneous Registers	123
2.19.1	MC_DIMM_CLK_RATIO_STATUS	123
2.19.2	MC_DIMM_CLK_RATIO	124
3	DIMM Population Requirements	125
3.1	General Population Requirements	125
3.2	Populating DIMMs Within a Channel	126
3.2.1	DIMM Population for Three Slots per Channel	126
3.2.2	DIMM Population for Two Slots per Channel	128

Tables

1-1	References	17
2-1	Functions Specifically Handled by the Processor	22
2-2	Device 0, Function 0: Generic Non-core Registers	23
2-3	Device 0, Function 1: System Address Decoder Registers	24
2-4	Device 2, Function 0: Intel QPI Link 0 Registers	25
2-5	Device 2, Function 1: Intel QPI Physical 0 Registers	26
2-6	Device 2, Function 4: Intel QPI Link 1 Registers1	27
2-7	Device 2, Function 5: Intel QPI Physical 1 Registers	28
2-8	Device 3, Function 0: Integrated Memory Controller Registers	29
2-9	Device 3, Function 1: Target Address Decoder Registers	30
2-10	Device 3, Function 2: Integrated Memory Controller RAS Registers	31
2-11	Device 3, Function 4: Integrated Memory Controller Test Registers	32
2-12	Device 4, Function 0: Integrated Memory Controller Channel 0 Control Registers	33
2-13	Device 4, Function 1: Integrated Memory Controller Channel 0 Address Registers	34
2-14	Device 4, Function 2: Integrated Memory Controller Channel 0 Rank Registers	35
2-15	Device 4, Function 3: Integrated Memory Controller Channel 0 Thermal Control Registers	36
2-16	Device 5, Function 0: Integrated Memory Controller Channel 1 Control Registers	37
2-17	Device 5, Function 1: Integrated Memory Controller Channel 1 Address Registers	38
2-18	Device 5, Function 2: Integrated Memory Controller Channel 1 Rank Registers	39
2-19	Device 5, Function 3: Integrated Memory Controller Channel 1 Thermal Control Registers	40
2-20	Device 6, Function 0: Integrated Memory Controller Channel 2 Control Registers	41
2-21	Device 6, Function 1: Integrated Memory Controller Channel 2 Address Registers	42
2-22	Device 6, Function 2: Integrated Memory Controller Channel 2 Rank Registers	43
2-23	Device 6, Function 3: Integrated Memory Controller Channel 2 Thermal Control Registers	44



3-1	Key Parameters for DIMM Configurations	125
3-2	RDIMM Population Configurations within a Channel for Three Slots per Channel	127
3-3	UDIMM Population Configurations within a Channel for Three Slots per Channel	127
3-4	MetaSDRAM* R-DIMM Population Configurations within a Channel for Three Slots per Channel	128
3-5	RDIMM Population Configurations Within a Channel for Two Slots per Channel	129
3-6	UDIMM Population Configurations within a Channel for Two Slots per Channel	129
3-7	MetaSDRAM R-DIMM Population Configurations within a Channel for Two Slots per Channel	129

Figures


3-1	DIMM Population within a Channel for Three Slots per Channel	126
3-2	DIMM Population Within a Channel for Two Slots per Channel	128




Revision History

Reference Number	Revision Number	Description	Date
321322	001	Public release	March 2009
321322	002	Added Chapter 3 "DIMM Population Requirements"	April 2009

§





1 Introduction

The Intel® Xeon® Processor 5500 Series is the first generation DP server/workstation processor to implement 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.

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

This document provides register documentation and functional description of major functional areas of the processor non-core design such as the memory controller and Intel QPI logic, and additional features pertinent to implementation and operation of the processor.

The Intel Xeon Processor 5500 Series are multi-core processors, based on 45 nm process technology. Processor features vary by SKU 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), and Intel® Hyper-Threading 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://developer.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 Land Grid Array (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://developer.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 5500 Series** — The 2S server/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 $t_1, t_2, t_n, \dots, t_k$ 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	Reference #	Notes
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual</i> <ul style="list-style-type: none">• Volume 1: Basic Architecture• Volume 2A: Instruction Set Reference, A-M• Volume 2B: Instruction Set Reference, N-Z• Volume 3A: System Programming Guide, Part 1• Volume 3B: Systems Programming Guide, Part 2	253665 253666 253667 253668 253669	1
<i>Intel® 64 and IA-32 Architectures Optimization Reference Manual</i>	248966	1
<i>Intel® Xeon® Processor 5500 Series Specification Update</i>	321324	1
<i>Intel® Xeon® Processor 5500 Series Datasheet, Volume 1</i>	321321	1

Notes:

1. Document is available publicly at <http://www.intel.com>.



Introduction



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*, as well as the PCI Express enhanced configuration mechanism as specified in the *PCI Express Base Specification*. All the registers are organized by bus, device, function, etc. as defined in the *PCI Express Base Specification*. 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 SKU, not all of the register descriptions in this document apply to all processors. This document highlights registers which do not apply to all processor SKUs. 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.
RW0C	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.
RW0S	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> 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 2C40h. 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 2, Function 4 contains the Intel® QuickPath configuration registers for Intel® QuickPath Interconnect Link 1 and resides at DID of 2C14h. Device 2, Function 5 contains the physical layer registers for Intel QPI Link 1 and resides at DID of 2C15h. Functions 4 and 5 only apply to processors with two Intel QPI links.
- **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	2C40h	0	0
	Intel® QuickPath Architecture System Address Decoder	2C01h		1
	Intel QPI Link 0	2C10h	2	0
	Intel QPI Physical 0	2C11		1
	Intel QPI Link 1	2C14h		4 ¹
	Intel QPI Physical 1	2C15h		5 ¹
	Integrated Memory Controller Registers	2C18h	3	0
	Integrated Memory Controller Target Address Decoder	2C19h		1
	Integrated Memory Controller RAS Registers	2C1Ah		2 ²
	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:

1. Applies only to processors with two Intel QPI links.
2. Applies only to processors supporting mirroring and scrubbing RAS features.



2.4 Detailed Configuration Space Maps

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

DID	VID	00h	DESIRED_CORES	80h
PCISTS	PCICMD	04h		84h
CCR	RID	08h	MEMLOCK_STATUS	88h
HDR		0Ch		8Ch
		10h	MC_CFG_CONTROL	90h
		14h		94h
		18h		98h
		1Ch		9Ch
		20h		A0h
		24h		A4h
		28h		A8h
SID	SVID	2Ch		ACH
		30h	POWER_CNTRL_ERR_STATUS	B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
MAXREQUEST_LC		40h	CURRENT_UCLK_RATIO	C0h
MAXREQUEST_LS		44h		C4h
MAXREQUEST_LL		48h		C8h
		4Ch		CCh
		50h	MIRROR_PORT_CTL	D0h
		54h		D4h
		58h		D8h
		5Ch		DCh
MAX_RTIDS		60h	MIP_PH_CTR_L0	E0h
		64h	MIP_PH_PRT_L0	E4h
		68h		E8h
		6Ch		ECh
		70h	MIP_PH_CTR_L1	F0h
		74h	MIP_PH_PRT_L1	F4h
		78h		F8h
		7Ch		FCh



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

DID	VID	00h	SAD_DRAM_RULE_0	80h
PCISTS	PCICMD	04h	SAD_DRAM_RULE_1	84h
CCR	RID	08h	SAD_DRAM_RULE_2	88h
HDR		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
SID	SVID	2Ch		ACh
		30h		B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
SAD_PAM0123		40h	SAD_INTERLEAVE_LIST_0	C0h
SAD_PAM456		44h	SAD_INTERLEAVE_LIST_1	C4h
SAD_HEN		48h	SAD_INTERLEAVE_LIST_2	C8h
SAD_SMRAM		4Ch	SAD_INTERLEAVE_LIST_3	CCh
SAD_PCIEXBAR		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	00h		80h
PCISTS	PCICMD	04h		84h
CCR	RID	08h		88h
HDR		0Ch		8Ch
		10h		90h
		14h		94h
		18h		98h
		1Ch		9Ch
		20h		A0h
		24h		A4h
		28h		A8h
SID	SVID	2Ch		ACH
		30h		B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
QPI_QPILCP_L0	QPI_RMT_QPILP0_STAT_L0	40h		C0h
	QPI_RMT_QPILP1_STAT_L0	44h		C4h
QPI_QPILCL_L0	QPI_RMT_QPILP2_STAT_L0	48h		C8h
	QPI_RMT_QPILP3_STAT_L0	4Ch		CCh
QPI_QPILS_L0		50h		D0h
		54h		D4h
QPI_DEF_RMT_VN_CREDITS_L0		58h		D8h
		5Ch		DCh
		60h		E0h
		64h		E4h
		68h		E8h
		6Ch		ECh
		70h		F0h
		74h		F4h
		78h		F8h
		7Ch		FCh



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

DID	VID	00h	QPI_0_PH_PIS	80h
PCISTS	PCICMD	04h		84h
CCR	RID	08h		88h
HDR		0Ch		8Ch
		10h		90h
		14h	QPI_0_PH_PTV	94h
		18h		98h
		1Ch	QPI_0_PH_LDC	9Ch
		20h		A0h
		24h	QPI_0_PH_PRT	A4h
		28h		A8h
SID	SVID	2Ch		ACH
		30h		B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
		40h		C0h
		44h		C4h
		48h		C8h
		4Ch		CCh
QPI_0_PLL_STATUS		50h	QPI_0_PH_PMRO	D0h
QPI_0_PLL_RATIO		54h		D4h
		58h		D8h
		5Ch		DCh
		60h	QPI_0_EP_SR	E0h
		64h		E4h
QPI_0_PH_CPR		68h		E8h
QPI_0_PH_CTR		6Ch		ECh
		70h		F0h
		74h	QPI_0_EP_MCTR	F4h
		78h		F8h
		7Ch		FCh




Table 2-6. Device 2, Function 4: Intel QPI Link 1 Registers¹

DID		VID		00h		80h			
PCISTS		PCICMD		04h		84h			
CCR			RID	08h		88h			
BIST	HDR					0Ch	8Ch		
						10h	90h		
						14h	94h		
						18h	98h		
						1Ch	9Ch		
						20h	A0h		
						24h	A4h		
						28h	A8h		
SID		SVID		2Ch		ACH			
				30h	B0h				
				34h	B4h				
				38h	B8h				
				3Ch	BCh				
				QPI_QPILCP_L1				40h	C0h
				44h	C4h				
QPI_QPILCL_L1				48h	C8h				
				4Ch	CCh				
QPI_QPILS_L1				50h	D0h				
				54h	D4h				
QPI_DEF_RMT_VN_CREDITS_L1				58h	D8h				
				5Ch	DCh				
				60h	E0h				
				64h	E4h				
				68h	E8h				
				6Ch	ECh				
				70h	F0h				
				74h	F4h				
				78h	F8h				
				7Ch	FCh				

Note:
1. Applies only to processors with two Intel QPI links.



Table 2-7. Device 2, Function 5: Intel QPI Physical 1 Registers

DID	VID	00h	QPI_1_PH_PIS	80h
PCISTS	PCICMD	04h		84h
CCR	RID	08h		88h
HDR		0Ch		8Ch
		10h		90h
		14h	QPI_1_PH_PTV	94h
		18h		98h
		1Ch	QPI_1_PH_LDC	9Ch
		20h		A0h
		24h	QPI_1_PH_PRT	A4h
		28h		A8h
SID	SVID	2Ch		ACH
		30h		B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
		40h		C0h
		44h		C4h
		48h		C8h
		4Ch		CCh
QPI_1_PLL_STATUS		50h	QPI_1_PH_PMR0	D0h
QPI_1_PLL_RATIO		54h		D4h
		58h		D8h
		5Ch		DCh
		60h	QPI_1_EP_SR	E0h
		64h		E4h
QPI_1_PH_CPR		68h		E8h
QPI_1_PH_CTR		6Ch		ECh
		70h		F0h
		74h	QPI_1_EP_MCTR	F4h
		78h		F8h
		7Ch		FCh




Table 2-8. Device 3, Function 0: Integrated Memory Controller Registers

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



Table 2-9. Device 3, Function 1: Target Address Decoder Registers

DID	VID	00h	TAD_DRAM_RULE_0	80h
PCISTS	PCICMD	04h	TAD_DRAM_RULE_1	84h
CCR	RID	08h	TAD_DRAM_RULE_2	88h
HDR		0Ch	TAD_DRAM_RULE_3	8Ch
		10h	TAD_DRAM_RULE_4	90h
		14h	TAD_DRAM_RULE_5	94h
		18h	TAD_DRAM_RULE_6	98h
		1Ch	TAD_DRAM_RULE_7	9Ch
		20h		A0h
		24h		A4h
		28h		A8h
SID	SVID	2Ch		ACH
		30h		B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
		40h	TAD_INTERLEAVE_LIST_0	C0h
		44h	TAD_INTERLEAVE_LIST_1	C4h
		48h	TAD_INTERLEAVE_LIST_2	C8h
		4Ch	TAD_INTERLEAVE_LIST_3	CCh
		50h	TAD_INTERLEAVE_LIST_4	D0h
		54h	TAD_INTERLEAVE_LIST_5	D4h
		58h	TAD_INTERLEAVE_LIST_6	D8h
		5Ch	TAD_INTERLEAVE_LIST_7	DCh
		60h		E0h
		64h		E4h
		68h		E8h
		6Ch		ECh
		70h		F0h
		74h		F4h
		78h		F8h
		7Ch		FCh

Table 2-10. Device 3, Function 2: Integrated Memory Controller RAS Registers¹

DID	VID	00h	MC_COR_ECC_CNT_0	80h
PCISTS	PCICMD	04h	MC_COR_ECC_CNT_1	84h
CCR	RID	08h	MC_COR_ECC_CNT_2	88h
HDR		0Ch	MC_COR_ECC_CNT_3	8Ch
		10h	MC_COR_ECC_CNT_4	90h
		14h	MC_COR_ECC_CNT_5	94h
		18h		98h
		1Ch		9Ch
		20h		A0h
		24h		A4h
		28h		A8h
SID	SVID	2Ch		ACh
		30h		B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
		40h		C0h
		44h		C4h
MC_SSRCONTROL		48h		C8h
MC_SCRUB_CONTROL		4Ch		CCh
MC_RAS_ENABLES		50h		D0h
MC_RAS_STATUS		54h		D4h
		58h		D8h
		5Ch		DCh
MC_SSRSTATUS		60h		E0h
		64h		E4h
		68h		E8h
		6Ch		ECh
		70h		F0h
		74h		F4h
		78h		F8h
		7Ch		FCh

Notes:

1. Applies only to processors supporting registered DIMMs.



Table 2-11. Device 3, Function 4: Integrated Memory Controller Test Registers

DID	VID	00h	MC_TEST_PH_PIS	80h
PCISTS	PCICMD	04h		84h
CCR	RID	08h		88h
HDR		0Ch		8Ch
		10h		90h
		14h		94h
		18h		98h
		1Ch		9Ch
		20h		A0h
		24h		A4h
		28h	MC_TEST_PAT_GCTR	A8h
SID	SVID	2Ch		ACH
		30h	MC_TEST_PAT_BA	B0h
		34h		B4h
		38h		B8h
		3Ch	MC_TEST_PAT_IS	BCh
		40h	MC_TEST_PAT_DCD	C0h
		44h		C4h
		48h		C8h
		4Ch		CCh
MC_DIMM_CLK_RATIO_STATUS		50h		D0h
MC_DIMM_CLK_RATIO		54h		D4h
		58h		D8h
		5Ch		DCh
MC_TEST_ERR_RCV1		60h		E0h
MC_TEST_ERR_RCVO		64h		E4h
		68h		E8h
MC_TEST_PH_CTR		6Ch		ECh
		70h		F0h
		74h		F4h
		78h		F8h
		7Ch		FCh


Table 2-12. Device 4, Function 0: Integrated Memory Controller Channel 0 Control Registers

DID	VID	00h	MC_CHANNEL_0_RANK_TIMING_A	80h
PCISTS	PCICMD	04h	MC_CHANNEL_0_RANK_TIMING_B	84h
CCR	RID	08h	MC_CHANNEL_0_BANK_TIMING	88h
HDR		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
SID	SVID	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_WAQ_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
MC_CHANNEL_0_DIMM_RESET_CMD		50h	MC_CHANNEL_0_EW_BGF_OFFSET_SETTINGS	D0h
MC_CHANNEL_0_DIMM_INIT_CMD		54h	MC_CHANNEL_0_ROUND_TRIP_LATENCY	D4h
MC_CHANNEL_0_DIMM_INIT_PARAMS		58h	MC_CHANNEL_0_PAGETABLE_PARAMS1	D8h
MC_CHANNEL_0_DIMM_INIT_STATUS		5Ch		DCh
MC_CHANNEL_0_DDR3CMD		60h	MC_TX_BG_CMD_DATA_RATIO_SETTING_CH0	E0h
		64h	MC_TX_BG_CMD_OFFSET_SETTINGS_CH0	E4h
MC_CHANNEL_0_REFRESH_THROTTLE_SUPPORT		68h	MC_TX_BG_DATA_OFFSET_SETTINGS_CH0	E8h
		6Ch		ECh
MC_CHANNEL_0_MRS_VALUE_0_1		70h	MC_CHANNEL_0_ADDR_MATCH	F0h
MC_CHANNEL_0_MRS_VALUE_2		74h		F4h
		78h	MC_CHANNEL_0_ECC_ERROR_MASK	F8h
MC_CHANNEL_0_RANK_PRESENT		7Ch	MC_CHANNEL_0_ECC_ERROR_INJECT	FCh



Table 2-13. Device 4, Function 1: Integrated Memory Controller Channel 0 Address Registers

DID	VID	00h	MC_SAG_CH0_0	80h
PCISTS	PCICMD	04h	MC_SAG_CH0_1	84h
CCR	RID	08h	MC_SAG_CH0_2	88h
HDR		0Ch	MC_SAG_CH0_3	8Ch
		10h	MC_SAG_CH0_4	90h
		14h	MC_SAG_CH0_5	94h
		18h	MC_SAG_CH0_6	98h
		1Ch	MC_SAG_CH0_7	9Ch
		20h		A0h
		24h		A4h
		28h		A8h
SID	SVID	2Ch		ACH
		30h		B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
		40h		C0h
		44h		C4h
MC_DOD_CH0_0		48h		C8h
MC_DOD_CH0_1		4Ch		CCh
MC_DOD_CH0_2		50h		D0h
		54h		D4h
		58h		D8h
		5Ch		DCh
		60h		E0h
		64h		E4h
		68h		E8h
		6Ch		ECh
		70h		F0h
		74h		F4h
		78h		F8h
		7Ch		FCh



Table 2-14. Device 4, Function 2: Integrated Memory Controller Channel 0 Rank Registers

DID	VID	00h	MC_RIR_WAY_CH0_0	80h
PCISTS	PCICMD	04h	MC_RIR_WAY_CH0_1	84h
CCR	RID	08h	MC_RIR_WAY_CH0_2	88h
HDR		0Ch	MC_RIR_WAY_CH0_3	8Ch
		10h	MC_RIR_WAY_CH0_4	90h
		14h	MC_RIR_WAY_CH0_5	94h
		18h	MC_RIR_WAY_CH0_6	98h
		1Ch	MC_RIR_WAY_CH0_7	9Ch
		20h	MC_RIR_WAY_CH0_8	A0h
		24h	MC_RIR_WAY_CH0_9	A4h
		28h	MC_RIR_WAY_CH0_10	A8h
SID	SVID	2Ch	MC_RIR_WAY_CH0_11	ACH
		30h	MC_RIR_WAY_CH0_12	B0h
		34h	MC_RIR_WAY_CH0_13	B4h
		38h	MC_RIR_WAY_CH0_14	B8h
		3Ch	MC_RIR_WAY_CH0_15	BCh
MC_RIR_LIMIT_CH0_0		40h	MC_RIR_WAY_CH0_16	C0h
MC_RIR_LIMIT_CH0_1		44h	MC_RIR_WAY_CH0_17	C4h
MC_RIR_LIMIT_CH0_2		48h	MC_RIR_WAY_CH0_18	C8h
MC_RIR_LIMIT_CH0_3		4Ch	MC_RIR_WAY_CH0_19	CCh
MC_RIR_LIMIT_CH0_4		50h	MC_RIR_WAY_CH0_20	D0h
MC_RIR_LIMIT_CH0_5		54h	MC_RIR_WAY_CH0_21	D4h
MC_RIR_LIMIT_CH0_6		58h	MC_RIR_WAY_CH0_22	D8h
MC_RIR_LIMIT_CH0_7		5Ch	MC_RIR_WAY_CH0_23	DCh
		60h	MC_RIR_WAY_CH0_24	E0h
		64h	MC_RIR_WAY_CH0_25	E4h
		68h	MC_RIR_WAY_CH0_26	E8h
		6Ch	MC_RIR_WAY_CH0_27	ECh
		70h	MC_RIR_WAY_CH0_28	F0h
		74h	MC_RIR_WAY_CH0_29	F4h
		78h	MC_RIR_WAY_CH0_30	F8h
		7Ch	MC_RIR_WAY_CH0_31	FCh



Table 2-15. Device 4, Function 3: Integrated Memory Controller Channel 0 Thermal Control Registers

DID	VID	00h	MC_COOLING_COEF0	80h
PCISTS	PCICMD	04h	MC_CLOSED_LOOP0	84h
CCR	RID	08h	MC_THROTTLE_OFFSET0	88h
HDR		0Ch		8Ch
		10h		90h
		14h		94h
		18h	MC_RANK_VIRTUAL_TEMPO	98h
		1Ch	MC_DDR_THERM_COMMAND0	9Ch
		20h		A0h
		24h	MC_DDR_THERM_STATUS0	A4h
		28h		A8h
SID	SVID	2Ch		ACH
		30h		B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
		40h		C0h
		44h		C4h
		48h		C8h
MC_THERMAL_CONTROLO		4Ch		CCh
MC_THERMAL_STATUS0		50h		D0h
MC_THERMAL_DEFEATURE0		54h		D4h
		58h		D8h
		5Ch		DCh
MC_THERMAL_PARAMS_A0		60h		E0h
MC_THERMAL_PARAMS_B0		64h		E4h
		68h		E8h
		6Ch		ECh
		70h		F0h
		74h		F4h
		78h		F8h
		7Ch		FCh


Table 2-16. Device 5, Function 0: Integrated Memory Controller Channel 1 Control Registers

DID	VID	00h	MC_CHANNEL_1_RANK_TIMING_A	80h
PCISTS	PCICMD	04h	MC_CHANNEL_1_RANK_TIMING_B	84h
CCR	RID	08h	MC_CHANNEL_1_BANK_TIMING	88h
HDR		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
SID	SVID	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_WAQ_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
MC_CHANNEL_1_DIMM_RESET_CMD		50h	MC_CHANNEL_1_EW_BGF_OFFSET_SETTINGS	D0h
MC_CHANNEL_1_DIMM_INIT_CMD		54h	MC_CHANNEL_1_ROUND_TRIP_LATENCY	D4h
MC_CHANNEL_1_DIMM_INIT_PARAMS		58h	MC_CHANNEL_1_PAGETABLE_PARAMS1	D8h
MC_CHANNEL_1_DIMM_INIT_STATUS		5Ch		DCh
MC_CHANNEL_1_DDR3CMD		60h	MC_TX_BG_CMD_DATA_RATIO_SETTING_CH1	E0h
		64h	MC_TX_BG_CMD_OFFSET_SETTINGS_CH1	E4h
MC_CHANNEL_1_REFRESH_THROTTLE_SUPPORT		68h	MC_TX_BG_DATA_OFFSET_SETTINGS_CH1	E8h
		6Ch		ECh
MC_CHANNEL_1_MRS_VALUE_0_1		70h	MC_CHANNEL_1_ADDR_MATCH	F0h
MC_CHANNEL_1_MRS_VALUE_2		74h		F4h
		78h	MC_CHANNEL_1_ECC_ERROR_MASK	F8h
MC_CHANNEL_1_RANK_PRESENT		7Ch	MC_CHANNEL_1_ECC_ERROR_INJECT	FCh



Table 2-17. Device 5, Function 1: Integrated Memory Controller Channel 1 Address Registers

DID	VID	00h	MC_SAG_CH1_0	80h
PCISTS	PCICMD	04h	MC_SAG_CH1_1	84h
CCR	RID	08h	MC_SAG_CH1_2	88h
HDR		0Ch	MC_SAG_CH1_3	8Ch
		10h	MC_SAG_CH1_4	90h
		14h	MC_SAG_CH1_5	94h
		18h	MC_SAG_CH1_6	98h
		1Ch	MC_SAG_CH1_7	9Ch
		20h		A0h
		24h		A4h
		28h		A8h
SID	SVID	2Ch		ACh
		30h		B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
		40h		C0h
		44h		C4h
MC_DOD_CH1_0		48h		C8h
MC_DOD_CH1_1		4Ch		CCh
MC_DOD_CH1_2		50h		D0h
		54h		D4h
		58h		D8h
		5Ch		DCh
		60h		E0h
		64h		E4h
		68h		E8h
		6Ch		ECh
		70h		F0h
		74h		F4h
		78h		F8h
		7Ch		FCh




Table 2-18. Device 5, Function 2: Integrated Memory Controller Channel 1 Rank Registers

DID	VID	00h	MC_RIR_WAY_CH1_0	80h
PCISTS	PCICMD	04h	MC_RIR_WAY_CH1_1	84h
CCR	RID	08h	MC_RIR_WAY_CH1_2	88h
HDR		0Ch	MC_RIR_WAY_CH1_3	8Ch
		10h	MC_RIR_WAY_CH1_4	90h
		14h	MC_RIR_WAY_CH1_5	94h
		18h	MC_RIR_WAY_CH1_6	98h
		1Ch	MC_RIR_WAY_CH1_7	9Ch
		20h	MC_RIR_WAY_CH1_8	A0h
		24h	MC_RIR_WAY_CH1_9	A4h
		28h	MC_RIR_WAY_CH1_10	A8h
		2Ch	MC_RIR_WAY_CH1_11	ACH
SID	SVID	30h	MC_RIR_WAY_CH1_12	B0h
		34h	MC_RIR_WAY_CH1_13	B4h
		38h	MC_RIR_WAY_CH1_14	B8h
		3Ch	MC_RIR_WAY_CH1_15	BCh
		40h	MC_RIR_WAY_CH1_16	C0h
MC_RIR_LIMIT_CH1_0		44h	MC_RIR_WAY_CH1_17	C4h
MC_RIR_LIMIT_CH1_1		48h	MC_RIR_WAY_CH1_18	C8h
MC_RIR_LIMIT_CH1_2		4Ch	MC_RIR_WAY_CH1_19	CCh
MC_RIR_LIMIT_CH1_3		50h	MC_RIR_WAY_CH1_20	D0h
MC_RIR_LIMIT_CH1_4		54h	MC_RIR_WAY_CH1_21	D4h
MC_RIR_LIMIT_CH1_5		58h	MC_RIR_WAY_CH1_22	D8h
MC_RIR_LIMIT_CH1_6		5Ch	MC_RIR_WAY_CH1_23	DCh
MC_RIR_LIMIT_CH1_7		60h	MC_RIR_WAY_CH1_24	E0h
		64h	MC_RIR_WAY_CH1_25	E4h
		68h	MC_RIR_WAY_CH1_26	E8h
		6Ch	MC_RIR_WAY_CH1_27	ECh
		70h	MC_RIR_WAY_CH1_28	F0h
		74h	MC_RIR_WAY_CH1_29	F4h
		78h	MC_RIR_WAY_CH1_30	F8h
		7Ch	MC_RIR_WAY_CH1_31	FCh



Table 2-19. Device 5, Function 3: Integrated Memory Controller Channel 1 Thermal Control Registers

DID	VID	00h	MC_COOLING_COEF1	80h
PCISTS	PCICMD	04h	MC_CLOSED_LOOP1	84h
CCR	RID	08h	MC_THROTTLE_OFFSET1	88h
HDR		0Ch		8Ch
		10h		90h
		14h		94h
		18h	MC_RANK_VIRTUAL_TEMP1	98h
		1Ch	MC_DDR_THERM_COMMAND1	9Ch
		20h		A0h
		24h	MC_DDR_THERM_STATUS1	A4h
		28h		A8h
SID	SVID	2Ch		ACh
		30h		B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
		40h		C0h
		44h		C4h
		48h		C8h
MC_THERMAL_CONTROL1		4Ch		CCh
MC_THERMAL_STATUS1		50h		D0h
MC_THERMAL_DEFEATURE1		54h		D4h
		58h		D8h
		5Ch		DCh
		60h		E0h
MC_THERMAL_PARAMS_A1		64h		E4h
MC_THERMAL_PARAMS_B1		68h		E8h
		6Ch		ECh
		70h		F0h
		74h		F4h
		78h		F8h
		7Ch		FCh


Table 2-20. Device 6, Function 0: Integrated Memory Controller Channel 2 Control Registers

DID	VID	00h	MC_CHANNEL_2_RANK_TIMING_A	80h
PCISTS	PCICMD	04h	MC_CHANNEL_2_RANK_TIMING_B	84h
CCR	RID	08h	MC_CHANNEL_2_BANK_TIMING	88h
HDR		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
SID	SVID	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_WAQ_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
MC_CHANNEL_2_DIMM_RESET_CMD		50h	MC_CHANNEL_2_EW_BGF_OFFSET_SETTINGS	D0h
MC_CHANNEL_2_DIMM_INIT_CMD		54h	MC_CHANNEL_2_ROUND_TRIP_LATENCY	D4h
MC_CHANNEL_2_DIMM_INIT_PARAMS		58h	MC_CHANNEL_2_PAGETABLE_PARAMS1	D8h
MC_CHANNEL_2_DIMM_INIT_STATUS		5Ch		DCh
MC_CHANNEL_2_DDR3CMD		60h	MC_TX_BG_CMD_DATA_RATIO_SETTING_CH2	E0h
		64h	MC_TX_BG_CMD_OFFSET_SETTINGS_CH2	E4h
MC_CHANNEL_2_REFRESH_THROTTLE_SUPPORT		68h	MC_TX_BG_DATA_OFFSET_SETTINGS_CH2	E8h
		6Ch		ECh
MC_CHANNEL_2_MRS_VALUE_0_1		70h	MC_CHANNEL_2_ADDR_MATCH	F0h
MC_CHANNEL_2_MRS_VALUE_2		74h		F4h
		78h	MC_CHANNEL_2_ECC_ERROR_MASK	F8h
MC_CHANNEL_2_RANK_PRESENT		7Ch	MC_CHANNEL_2_ECC_ERROR_INJECT	FCh



Table 2-21. Device 6, Function 1: Integrated Memory Controller Channel 2 Address Registers

DID	VID	00h	MC_SAG_CH2_0	80h
PCISTS	PCICMD	04h	MC_SAG_CH2_1	84h
CCR	RID	08h	MC_SAG_CH2_2	88h
HDR		0Ch	MC_SAG_CH2_3	8Ch
		10h	MC_SAG_CH2_4	90h
		14h	MC_SAG_CH2_5	94h
		18h	MC_SAG_CH2_6	98h
		1Ch	MC_SAG_CH2_7	9Ch
		20h		A0h
		24h		A4h
		28h		A8h
SID	SVID	2Ch		ACh
		30h		B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
		40h		C0h
		44h		C4h
MC_DOD_CH2_0		48h		C8h
MC_DOD_CH2_1		4Ch		CCh
MC_DOD_CH2_2		50h		D0h
		54h		D4h
		58h		D8h
		5Ch		DCh
		60h		E0h
		64h		E4h
		68h		E8h
		6Ch		ECh
		70h		F0h
		74h		F4h
		78h		F8h
		7Ch		FCh



Table 2-22. Device 6, Function 2: Integrated Memory Controller Channel 2 Rank Registers

DID	VID	00h	MC_RIR_WAY_CH2_0	80h
PCISTS	PCICMD	04h	MC_RIR_WAY_CH2_1	84h
CCR	RID	08h	MC_RIR_WAY_CH2_2	88h
HDR		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
SID	SVID	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
MC_RIR_LIMIT_CH2_0		40h	MC_RIR_WAY_CH2_16	C0h
MC_RIR_LIMIT_CH2_1		44h	MC_RIR_WAY_CH2_17	C4h
MC_RIR_LIMIT_CH2_2		48h	MC_RIR_WAY_CH2_18	C8h
MC_RIR_LIMIT_CH2_3		4Ch	MC_RIR_WAY_CH2_19	CCh
MC_RIR_LIMIT_CH2_4		50h	MC_RIR_WAY_CH2_20	D0h
MC_RIR_LIMIT_CH2_5		54h	MC_RIR_WAY_CH2_21	D4h
MC_RIR_LIMIT_CH2_6		58h	MC_RIR_WAY_CH2_22	D8h
MC_RIR_LIMIT_CH2_7		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 2-23. Device 6, Function 3: Integrated Memory Controller Channel 2 Thermal Control Registers

DID	VID	00h	MC_COOLING_COEF2	80h
PCISTS	PCICMD	04h	MC_CLOSED_LOOP2	84h
CCR	RID	08h	MC_THROTTLE_OFFSET2	88h
HDR		0Ch		8Ch
		10h		90h
		14h		94h
		18h	MC_RANK_VIRTUAL_TEMP2	98h
		1Ch	MC_DDR_THERM_COMMAND2	9Ch
		20h		A0h
		24h	MC_DDR_THERM_STATUS2	A4h
		28h		A8h
SID	SVID	2Ch		ACH
		30h		B0h
		34h		B4h
		38h		B8h
		3Ch		BCh
		40h		C0h
		44h		C4h
		48h		C8h
MC_THERMAL_CONTROL2		4Ch		CCh
MC_THERMAL_STATUS2		50h		D0h
MC_THERMAL_DEFEATURE2		54h		D4h
		58h		D8h
		5Ch		DCh
		60h		E0h
MC_THERMAL_PARAMS_A2		64h		E4h
MC_THERMAL_PARAMS_B2		68h		E8h
		6Ch		ECh
		70h		F0h
		74h		F4h
		78h		F8h
		7Ch		FCh



2.5 PCI Standard Registers

These registers appear in every function for every device.

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 0: A Stepping 1: A Stepping 2: B Stepping 4: C Stepping 5: D Stepping Others: RSVD

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 PCICMD - 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	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. 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. 1: Legacy Interrupt mode is disabled 0: Legacy Interrupt mode is enabled
9	RO	0	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.
8	RO	0	SERRE: SERR Message Enable This bit is a global enable bit for this devices SERR messaging. This host bridge will not implement SERR messaging. This bit is hardwired to 0. Writes to this bit position have no effect. If SERR is used for error generation, then this bit must be RW and enable/disable SERR signaling.
7	RO	0	IDSELWCC: IDSEL Stepping/Wait Cycle Control Per PCI 2.3 spec this bit is hardwired to 0. Writes to this bit position have no effect.
6	RO	0	PERRE: Parity Error Response Enable Parity error is not implemented in this host bridge. This bit is hardwired to "0". Writes to this bit position have no effect.
5	RO	0	VGAPSE: VGA palette snoop Enable This host bridge does not implement this bit. This bit is hardwired to a "0". Writes to this bit position have no effect.
4	RO	0	MWIEN: Memory Write and Invalidate Enable This host bridge will never issue memory write and invalidate commands. This bit is therefore hardwired to "0". Writers to this bit position will have no effect.
3	RO	0	SCE: Special Cycle Enable This host bridge does not implement this bit. This bit is hardwired to a "0". Writers to this bit position will have no effect.
2	RO	1	BME: Bus Master Enable This host bridge is always enabled as a master. This bit is hardwired to a "1". Writes to this bit position have no effect.
1	RO	1	MSE: Memory Space Enable This host bridge always allows access to main memory. This bit is not implemented and is hardwired to "1". Writes to this bit position have no effect.
0	RO	0	IOAE: Access Enable This bit is not implemented in this host bridge and is hardwired to "0". Writes to this bit position have no effect.



2.5.8 PCI STS - 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". Writes to this bit position have no effect.
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. Writes to this bit position have no effect.
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. Writes to this bit position have no effect.
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. Writes to this bit position have no effect.
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". Writes to this bit position have no effect.
10:9	RO	0	DEVSEL Timing (DEVT) These bits are hardwired to "00". Writes to these bit positions have no effect. 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". Writes to this bit position have no effect.
7	RO	1	Fast Back-to-Back (FB2B) This bit is hardwired to "1". Writes to this bit position have no effect. 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. Must be 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	Capability List (CLIST) 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. If no capability structures are implemented, this bit is hardwired to 0.
3	RO	0	Interrupt Status: 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. If this device does not generate interrupts, then this bit is not implemented (RO and reads returns 0).
2:0	RO	0	Reserved

2.6 Generic Non-core Registers

2.6.1 MAXREQUEST_LC

Maximum requests expected from the chipset (number of TAD home trackers allocated to chipset). The maximum RTID value that may be used is one less than this number. Home trackers are allocated in groups of 8, so bits 2:0 of the register may not be written, and bits 5:3 indicate how many groups of 8 are allocated.

Device: 0 Function: 0 Offset: 40h Access as a Dword			
Bit	Type	Reset Value	Description
5:3	RW	3	VALUE. Maximum TAD requests from chipset (allocated in groups of 8).



2.6.2 MAXREQUEST_LS

Maximum requests expected from the sibling (number of TAD home trackers allocated to sibling). The maximum RTID value that may be used is one less than this number. Home Trackers are allocated in groups of 8, so bits 2:0 of the register may not be written, and bits 5:3 indicate how many groups of 8 are allocated.

Device: 0 Function: 0 Offset: 44h Access as a Dword			
Bit	Type	Reset Value	Description
5:3	RW	2	VALUE. Maximum TAD requests from sibling (allocated in groups of 8).

2.6.3 MAXREQUEST_LL

Maximum requests expected from local accesses (number of TAD home trackers allocated to the local queue). The maximum RTID value that may be used is one less than this number. Home Trackers are allocated in groups of 8, so bits 2:0 of the register may not be written, and bits 5:3 indicate how many groups of 8 are allocated.

Device: 0 Function: 0 Offset: 48h Access as a Dword			
Bit	Type	Reset Value	Description
5:3	RW	3	VALUE. Maximum TAD requests from local accesses (allocated in groups of 8).

2.6.4 MAX_RTIDS

Maximum number of RTIDs other homes have. How many requests can this caching agent send to the other home agents. This number is one more than the highest numbered RTID to use. Note these values reset to 2, and need to be increased by BIOS to whatever the home agents can support.

Device: 0 Function: 0 Offset: 60h Access as a Dword			
Bit	Type	Reset Value	Description
21:16	RW	2	LOCAL_MC. Maximum number of RTIDs for the local home agent.
13:8	RW	2	SIBLING. Maximum number of RTIDs for the sibling home agent.
5:0	RW	2	CHIPSET. Maximum number of RTIDs for the IOH home agent.



2.6.5 DESIRED_CORES

Number of cores, threads BIOS wants to exist on the next reset. A processor reset must be used for this register to take affect. Note programing this register to a value higher than the product has cores, should not be done. Which cores are removed is not defined and is implementation dependent. This does not result in all of the power savings of a reduced number of core product, but does save more power than even the deepest sleep state.

Device: 0 Function: 0 Offset: 80h Access as a Dword			
Bit	Type	Reset Value	Description
16	RW1S	0	LOCK. Once written to 1, changes to this register cannot be made.
8	RWL	0	MT_DISABLE. Disables multi-threading (2 logical threads per core) in all cores if set to 1.
1:0	RWL	0	CORE_COUNT. 00: max number (default value) 01 - 1 core 10 - 2 cores

2.6.6 MEMLOCK_STATUS

Status register for various Memory and Control Register functions that can be locked down.

Device: 0 Function: 0 Offset: 88h Access as a Dword			
Bit	Type	Reset Value	Description
9	RO	-	MEM_LOCKED_REMOTE. Any access to local memory from another agent (i.e. everybody but this processor) is aborted. Can only be unlocked when in Authenticated Code Mode.
8	RO	-	MEM_LOCKED_LOCAL. Any Access to local memory from this processor is aborted. Can only be unlocked when in Authenticated Code Mode.
1	RO	-	MEM_CFG_USER_LOCKED. Locks same as MEM_CFG_LOCKED but user controlled lockable by MC_CFG_CONTROL; unlockable via MC_CFG_CONTROL csr(0x0090).
0	RO	-	MEM_CFG_LOCKED. All Configuration registers dealing with memory and address programming are locked down and cannot be changed. This includes all registers in Device 3 Function [0,1], Device 4,5,6 Function 0, Device 4,5,6 Function 1, Device 4,5,6 Function 2, and most registers in Device 0 Function 1. But does not include the memory controller thermal registers, or SAD_PAM0123, SAD_PAM456, SAD_SMRAM registers.



2.6.7 MC_CFG_CONTROL

This register locks and unlocks write access to the Uncore configuration. BIOS must write a "1" to the MC_CFG_LOCK bit after reset to allow the Integrated Memory Controller to start accepting requests. It may subsequently be unlocked by writing a "1" to the MC_CFG_UNLOCK bit and a "0" to the MC_CFG_LOCK bit without affecting memory traffic.

Device: 0 Function: 0 Offset: 90h Access as a Dword			
Bit	Type	Reset Value	Description
1	WO	0	MC_CFG_UNLOCK. Unlocks Integrated Memory Controller configuration registers without CPU reset. This bit does NOT unlock any other lock type without a CPU reset.
0	WO	0	MC_CFG_LOCK. Locks Integrated Memory Controller configuration registers. Writes are no longer allowed to the configuration registers.

2.6.8 POWER_CNTRL_ERR_STATUS

Power management Error Status register.

Device: 0 Function: 0 Offset: B0h Access as a Qword			
Bit	Type	Reset Value	Description
63	RO	-	VAL. MC7_STATUS Register Valid. Indicates if the register is valid. 0: Not Valid 1: Valid
62	RO	-	OVER. Machine Check Overflow Flag. Indicates (when set) that a machine-check error occurred while the results of a previous error were still in the error-reporting register bank (that is, the VAL bit was already set in the IA32_MC7_STATUS register). The processor sets the OVER flag and software is responsible for clearing it. In general, enabled errors are written over disabled errors, and uncorrected errors are written over corrected errors. Uncorrected errors are not written over previous valid uncorrected errors. 0: No Overflow 1: Overflow
61	RO	-	UC. Error Uncorrected Flag. Indicates (when set) that the processor did not or was not able to correct the error condition. When cleared, this flag indicates that the processor was able to correct the error condition. 0: Corrected 1: Uncorrected
60	RO	-	EN. Error Enabled Flag. Indicates (when set) that the error was enabled by the associated EEj bit of the IA32_MC7_CTL register. 0: Not Enabled 1: Enabled



Device: 0 Function: 0 Offset: B0h Access as a Qword			
59	RO	-	MISCV. IA32_MC7_MISC. Register Valid Flag. Indicates (when set) that the IA32_MC7_MISC register contains additional information regarding the error. When clear, this flag indicates that the IA32_MC7_MISC register is either not implemented or does not contain additional information regarding the error. Do not read these registers if they are not implemented in the processor.
58	RO	-	ADDRV. IA32_MC7_ADDR. Register Valid Flag. Indicates (when set) that the IA32_MC7_ADDR register contains the address where the error occurred. When clear, this flag indicates that the IA32_MC7_ADDR register is either not implemented or does not contain the address where the error occurred. Do not read these registers if they are not implemented in the processor.
57	RO	-	PCC. Processor context corrupt flag. Indicates (when set) that the state of the processor might have been corrupted by the error condition detected and that reliable restarting of the processor may not be possible. When cleared, this flag indicates that the error did not affect the processor's state. 0: Not Corrupt 1: Corrupt
56:32	-	-	RSVD.
31:16	RO	-	MODEL SPECIFIC ERROR CODE. Specifies the model specific error code that uniquely identifies the machine-check error condition detected. The following list describes the error codes that may be found on the processor. 0x0000: No Error 0x0300: Unexpected reset error. Processor boot failed. 0x0800: PMReq or CmpD received was illegal in the current context. 0x0A00: Illegal PMReq request detected under S3, S4 or S5. 0x0D00: Invalid S-state transition requested. 0x1100: Platform / CPU VID controller mismatch. Processor boot failed. 0x1A00: Platform / CPU MSID mismatch. Processor boot failed. 0x2000: QPI training error.
15:0	RO	-	MCA ERROR CODE FIELD. Specifies the machine-check architecture-defined error code for the machine-check error condition detected. The machine-check architecture-defined error codes are guaranteed to be the same for all IA-32 processors that implement the machine-check architecture. See Section 14.7 of the Software Developers Manual, Vol 3A, "Interpreting the MCA Error Codes," and Appendix E, "Interpreting Machine-Check Error Codes", for information on machine-check error codes.

2.6.9 CURRENT_UCLK_RATIO

Status Register reporting the current Uncore Clk Ratio relative to BCLK (133Mhz). This is the clock in which the Last Level Cache (LLC) runs.

Device: 0 Function: 0 Offset: C0h Access as a Dword			
Bit	Type	Reset Value	Description
15	RW	0	RSVD.
14:8	RW	12	RSVD.
6:0	RO	-	UCLK. The current UCLK ratio



2.6.10 MIRROR_PORT_CTL

Mirror Port physical layer control register.

Device: 0 Function: 0 Offset: D0h Access as a Dword			
Bit	Type	Reset Value	Description
7	RW	0	SPARE. Spare MiP control register bits.
6	RW	0	DSBL_ENH_MPRX_SYNC. When set, it disables the enhancing synchronization scheme for the MiP_Rx.
5	RW	0	MIP_GO_10. When set, the Mip_Tx and Mip_Rx go to L0 directly from Config_FlitLock.
4	RW	0	MIP_RX_CRC_SQUASH. When set, replaces CRC errors with CRC special packet on MiP Rx.
3	RW	0	MIP_RX_PORT_SEL. Port select for MiP Rx. _PORT_SEL0=QPI Port 0. _PORT_SEL1=QPI Port 1.
2	RW	0	MIP_TX_PORT_SEL. Port select for MiP Tx. _PORT_SEL0=QPI Port 0. _PORT_SEL1=QPI Port 1.
1	RW	1	MIP_RX_ENABLE. Enables the Rx portion of the mirror port.
0	RW	1	MIP_TX_ENABLE. Enables the Tx portion of the mirror port.

2.6.11 MIP_PH_CTR_L0 MIP_PH_CTR_L1

Mirror Port Physical Layer Control Register.

Device: 0 Function: 0 Offset: E0h, F0h Access as a Dword			
Bit	Type	Reset Value	Description
27	RW	0	LA_LOAD_DISABLE. Disables the loading of the effective values of the Intel® QuickPath CSRs when set.
23	RW	0	ENABLE_PRBS. Enables LFSR pattern during bitlock/training.
22	RW	0	ENABLE_SCRAMBLE. Enables data scrambling through LFSR.
14	RW	1	DETERMINISM_MODE. Sets determinism mode of operation.
13	RW	1	DISABLE_AUTO_COMP. Disables automatic entry into compliance.
12	RW	0	INIT_FREEZE. When set, freezes the FSM when initialization aborts.
10:8	RW	0	INIT_MODE. Initialization mode that determines altered initialization modes.
7	RW	0	LINK_SPEED. Identifies slow speed or at-speed operation for the Intel QPI port.
5	RW	1	PHYINITBEGIN. Instructs the port to start initialization.
4	RW	0	SINGLE_STEP. Enables single step mode.
3	RW	0	LAT_FIX_CTL. If set, instructs the remote agent to fix the latency.
2	RW	0	BYPASS_CALIBRATION. Indicates the physical layer to bypass calibration.



Device: 0 Function: 0 Offset: E0h, F0h Access as a Dword			
Bit	Type	Reset Value	Description
1	RW	0	RESET_MODIFIER. Modifies soft reset to default reset when set.
0	RW1S	0	PHY_RESET. Physical Layer Reset. Note while this register is locked after going to FAST speed L0, this bit is not locked.

2.6.12 MIP_PH_PRT_L0 MIP_PH_PRT_L1

Mirror Port periodic retraining timing register.

Device: 0 Function: 0 Offset: E4h, F4h Access as a Dword			
Bit	Type	Reset Value	Description
21:16	RW	29	RETRAIN_PKT_CNT. Retraining packet count.
13:10	RW	11	EXP_RETRAIN_INTERVAL. Exponential count for retraining interval.
7:0	RW	3	RETRAIN_INTERVAL. Periodic retraining interval. A value of 0 indicates retraining is disabled.


2.7 SAD - System Address Decoder Registers

2.7.1 SAD_PAM0123

Register for legacy dev0func0 90h-93h address space.

Device: 0 Function: 1 Offset: 40h Access as a Dword			
Bit	Type	Reset Value	Description
29:28	RW	0	PAM3_HIENABLE. 0D4000-0D7FFF Attribute (HIENABLE) This field controls the steering of read and write cycles that address the BIOS area from 0D4000 to 0D7FFF. 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.
25:24	RW	0	PAM3_LOENABLE. 0D0000-0D3FFF Attribute (LOENABLE) This field controls the steering of read and write cycles that address the BIOS area from 0D0000 to 0D3FFF 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.

Register Description



Device: 0 Function: 1 Offset: 40h Access as a Dword			
Bit	Type	Reset Value	Description
21:20	RW	0	PAM2_HIENABLE. 0CC000-0CFFFF Attribute (HIENABLE) This field controls the steering of read and write cycles that address the BIOS area from 0CC000 to 0CFFFF. 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 sent 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. 0C8000-0CBFFF Attribute (LOENABLE) This field controls the steering of read and write cycles that address the BIOS area from 0C8000 to 0CBFFF. 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 sent 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. 0C4000-0C7FFF Attribute (HIENABLE) This field controls the steering of read and write cycles that address the BIOS area from 0C4000 to 0C7FFF. 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 sent 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. 0C0000-0C3FFF Attribute (LOENABLE) This field controls the steering of read and write cycles that address the BIOS area from 0C0000 to 0C3FFF. 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 sent to DRAM. Reads are serviced by ESI. 11: Normal DRAM Operation: All reads and writes are serviced by DRAM.
5:4	RW	0	PAMO_HIENABLE. 0F0000-0FFFFF Attribute (HIENABLE) This field controls the steering of read and write cycles that address the BIOS area from 0F0000 to 0FFFFF. 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 sent to DRAM. Reads are serviced by ESI. 11: Normal DRAM Operation: All reads and writes are serviced by DRAM.



2.7.2 SAD_PAM456

Register for legacy dev0func0 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. 0EC000-0EFFFF Attribute (HIENABLE) This field controls the steering of read and write cycles that address the BIOS area from 0EC000 to 0EFFFF. 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. 0E4000-0E7FFF Attribute (HIENABLE) This field controls the steering of read and write cycles that address the BIOS area from 0E4000 to 0E7FFF. 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. 0E0000-0E3FFF Attribute (LOENABLE) This field controls the steering of read and write cycles that address the BIOS area from 0E0000 to 0E3FFF. 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. 0DC000-0DFFFF Attribute (HIENABLE) This field controls the steering of read and write cycles that address the BIOS area from 0DC000 to 0DFFFF. 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. 0D8000-0DBFFF Attribute (LOENABLE) This field controls the steering of read and write cycles that address the BIOS area from 0D8000 to 0DBFFF. 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.



2.7.3 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. 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 16MB.

2.7.4 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, PCIEXBAR, (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 SMRAM Enable (G_SMRAME). If set to a 1, then Compatible SMRAM functions are enabled, providing 128 KB of DRAM accessible at the A0000h address while in SMM (ADSB with SMM decode). To enable Extended SMRAM 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 A0000 and BFFFF is supported so this field is hardwired to 010.



2.7.5 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: 256MB. 001: Reserved. 010: Reserved. 011: Reserved. 100: Reserved. 101: Reserved. 110: 64MB. 111: 128MB.
0	RW	0	ENABLE. Enable for PCIEBAR address space. Editing size should not be done without also enabling range.

2.7.6 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

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.



2.7.7 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

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.
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.

2.8 Intel QPI Link Registers

2.8.1 QPI_QPILCP_L0 QPI_QPILCP_L1

Intel QPI Link Capability. Function 4 in the below table applies only to processors with two Intel QPI links.

Device: 2 Function: 0, 4 Offset: 40h Access as a Dword			
Bit	Type	Reset Value	Description
27:26	RO	-	VNO_CRDTS_DATA. VNO Credits per Data MC 00 - 0 credits 01 - 1 10 - 2 to 8 11 - RSVD



Device: 2 Function: 0, 4 Offset: 40h Access as a Dword			
Bit	Type	Reset Value	Description
23:22	RO	-	VNO_CRDTS_NDATA. VNO Credits per Non-Data MC 00 - 0 credits 01 - 1 10 - 2 to 8 11 - RSVD
21:16	RO	-	VNA_CRDTS. VNA Credits / 8, after rounding down.
11	RO	-	CRC_SUPPORT. CRC Mode Support. 0 - 8b CRC. 1 - RSVD
9:8	RO	-	FLIT_INTERLEAVE. Flit Interleave. 00 - Idle/Null flit only. 01 - Command Insert Interleave. 10 - RSVD. 11 - RSVD.
7:0	RO	-	QPI_VER. Intel QPI Version Number 0 - Rev 1.0 !0 - RSVD.

2.8.2 QPI_QPI_LCL_L0 QPI_QPI_LCL_L1

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 L0s mode at the transmitter. This bit should be ANDed with the receive L0s capability bit received during parameter exchange to determine if a transmitter is allowed to enter into L0s. This is NOT a bit that determines the capability of a device.



2.8.3 QPI_QPILS_L0 QPI_QPILS_L1

Intel QPI Link Status.

Device: 2 Function: 0, 4 Offset: 50h Access as a Dword			
Bit	Type	Reset Value	Description
31	RO	-	CHIPSET_LINK . Indicates that the local physical link is connected to the IOH.

2.8.4 QPI_DEF_RMT_VN_CREDITS_L0 QPI_DEF_RMT_VN_CREDITS_L1

This is the control register that houses the default values of available remote credits to be transmitted to the remote agent for the remote Tx use.

Device: 2 Function: 0, 4 Offset: 58h Access as a Dword			
Bit	Type	Reset Value	Description
18:12	RW	100	VNA . VNA Credits.
11:10	RW	1	NCS . NCS Channel VNO Credits.
9:8	RW	1	NCB . NCB Channel VNO Credits.
7:6	RW	1	DRS . DRS Channel VNO Credits.
5:4	RW	1	NDR . NDRChannel VNO Credits.
3:2	RW	1	SNP . SNP Channel VNO Credits.
1:0	RW	1	HOM . HOMChannel VNO Credits.

2.8.5 QPI_RMT_QPILPO_STAT_L0 QPI_RMT_QPILPO_STAT_L1

Remote's QPI Parameter 0 Value register.

Device: 2 Function: 0, 4 Offset: C0h Access as a Dword			
Bit	Type	Reset Value	Description
23:16	RO	-	LLR_WRAP_VALUE . Value after which the LLR sequence counter wraps.
14:8	RO	-	NodeID_OFFSET . Node ID offset for the sending agent.
7:5	RO	-	NodeID . Number of Node IDs of the transmitting agent.
4:0	RO	-	PORT_NUMBER . Sender's port number.



2.8.6 QPI_RMT_QPILP1_STAT_L0 QPI_RMT_QPILP1_STAT_L1


Remote's QPI Parameter 1 Value register.

Device: 2 Function: 0, 4 Offset: C4h Access as a Dword			
Bit	Type	Reset Value	Description
9	RO	-	L1_SUPPORT. Indicates the remote agent's ability to support L1 state.
8	RO	-	LOP_SUPPORT. Indicates the remote agent's ability to support LOP state.
7	RO	-	LOS_SUPPORT. Indicates the remote agent's ability to support LOS state.
6	RO	-	RX_CII_SUPPORT. Indicates the remote agent's ability to receive CII data.
5	RO	-	PREFERRED_TX_SDI_MODE. Indicates the ability of the remote agent transmitter to send scheduled data interleave data.
4	RO	-	RCV_SDI_SUPPORT. Indicates remote agent can receive scheduled data interleave data.
3:2	RO	-	PREFERRED_TX_CRC_MODE. Preferred send mode for the remote transmitter. 00: No CRC 01: 8b CRC 10: 16b rolling CRC 11: RSVD
1:0	RO	-	RCV_CRC_MODE_SUPPORTED. CRC modes that the remote agent supports. 00: RSVD 01: 8b CRC 10: 16b and 8b CRC 11: RSVD

2.8.7 QPI_RMT_QPILP2_STAT_L0 QPI_RMT_QPILP2_STAT_L1

Remote's QPI Parameter 2 Value register.

Device: 2 Function: 0, 4 Offset: C8h Access as a Dword			
Bit	Type	Reset Value	Description
31	RO	-	Agent_000_Caching. Indicates agent 000 is a caching agent.
30	RO	-	Agent_000_Home. Indicates agent 000 is a home agent.
29	RO	-	Agent_000_IO_Proxy. Indicates agent 000 is an IO Proxy agent.
28	RO	-	RSVD.
26	RO	-	Agent_000_Router. Indicates agent 000 is a router agent.
25	RO	-	Agent_000_Firmware. Indicates agent 000 is a firmware agent.
24	RO	-	Agent_000_Config. Indicates agent 000 is a configuration agent.
23	RO	-	Agent_001_Caching. Indicates agent 001 is a caching agent.
22	RO	-	Agent_001_Home. Indicates agent 001 is a home agent.
21	RO	-	Agent_001_IO_Proxy. Indicates agent 001 is an IO Proxy agent.



Device: 2 Function: 0, 4 Offset: C8h Access as a Dword			
Bit	Type	Reset Value	Description
20	RO	-	RSVD.
18	RO	-	Agent_001_Router. Indicates agent 001 is a router agent.
17	RO	-	Agent_001_Firmware. Indicates agent 001 is a firmware agent.
16	RO	-	Agent_001_Config. Indicates agent 001 is a configuration agent.
15	RO	-	Agent_010_Caching. Indicates agent 010 is a caching agent.
14	RO	-	Agent_010_Home. Indicates agent 010 is a home agent.
13	RO	-	Agent_010_IO_Proxy. Indicates agent 010 is an IO Proxy agent.
12	RO	-	RSVD.
10	RO	-	Agent_010_Router. Indicates agent 010 is a router agent.
9	RO	-	Agent_010_Firmware. Indicates agent 010 is a firmware agent.
8	RO	-	Agent_010_Config. Indicates agent 010 is a configuration agent.
7	RO	-	Agent_011_Caching. Indicates agent 011 is a caching agent.
6	RO	-	Agent_011_Home. Indicates agent 011 is a home agent.
5	RO	-	Agent_011_IO_Proxy. Indicates agent 011 is an IO Proxy agent.
4	RO	-	RSVD.
2	RO	-	Agent_011_Router. Indicates agent 011 is a router agent.
1	RO	-	Agent_011_Firmware. Indicates agent 011 is a firmware agent.
0	RO	-	Agent_011_Config. Indicates agent 011 is a configuration agent.

2.8.8 QPI_RMT_QPILP3_STAT_LO QPI_RMT_QPILP3_STAT_L1

Remote's QPI Parameter 3 Value register.

Device: 2 Function: 0, 4 Offset: CCh Access as a Dword			
Bit	Type	Reset Value	Description
31	RO	-	Agent_100_Caching. Indicates agent 100 is a caching agent.
30	RO	-	Agent_100_Home. Indicates agent 100 is a home agent.
29	RO	-	Agent_100_IO_Proxy. Indicates agent 100 is an IO Proxy agent.
28	RO	-	RSVD.
26	RO	-	Agent_100_Router. Indicates agent 100 is a router agent.
25	RO	-	Agent_100_Firmware. Indicates agent 100 is a firmware agent.
24	RO	-	Agent_100_Config. Indicates agent 100 is a configuration agent.
23	RO	-	Agent_101_Caching. Indicates agent 101 is a caching agent.
22	RO	-	Agent_101_Home. Indicates agent 101 is a home agent.
21	RO	-	Agent_101_IO_Proxy. Indicates agent 101 is an IO Proxy agent.
20	RO	-	RSVD.



Device: 2 Function: 0, 4 Offset: CCh Access as a Dword			
Bit	Type	Reset Value	Description
18	RO	-	Agent_101_Router. Indicates agent 101 is a router agent.
17	RO	-	Agent_101_Firmware. Indicates agent 101 is a firmware agent.
16	RO	-	Agent_101_Config. Indicates agent 101 is a configuration agent.
15	RO	-	Agent_110_Caching. Indicates agent 110 is a caching agent.
14	RO	-	Agent_110_Home. Indicates agent 110 is a home agent.
13	RO	-	Agent_110_IO_Proxy. Indicates agent 110 is an IO Proxy agent.
12	RO	-	RSVD.
10	RO	-	Agent_110_Router. Indicates agent 110 is a router agent.
9	RO	-	Agent_110_Firmware. Indicates agent 110 is a firmware agent.
8	RO	-	Agent_110_Config. Indicates agent 110 is a configuration agent.
7	RO	-	Agent_111_Caching. Indicates agent 111 is a caching agent.
6	RO	-	Agent_111_Home. Indicates agent 111 is a home agent.
5	RO	-	Agent_111_IO_Proxy. Indicates agent 111 is an IO Proxy agent.
4	RO	-	RSVD.
2	RO	-	Agent_111_Router. Indicates agent 111 is a router agent.
1	RO	-	Agent_111_Firmware. Indicates agent 111 is a firmware agent.
0	RO	-	Agent_111_Config. Indicates agent 111 is a configuration agent.

2.9 Intel QPI Physical Layer Registers

2.9.1 QPI_0_PH_CPR QPI_1_PH_CPR

Intel QPI Physical Layer Capability Register.

Device: 2 Function: 1, 5 Offset: 68h Access as a Dword			
Bit	Type	Reset Value	Description
29	RO	-	LFSR_POLYNOMIAL. Agent's ITU polynomial capability for loopback.
28:24	RO	-	NUMBER_OF_TX_LANES. Number of Tx lanes with which an implementation can operate for full width. Bit 28 - If set, 20 lanes. The bit indicating the maximum lanes will determine the number of control/status bits implemented in Tx/Rx Data lane Control/Status Registers.
23	RO	-	PRBS_CAPABILITY. If set, implementation is capable of using specified pattern in bitlock/retraining.
22	RO	-	SCRAMBLE_CAPABILITY. If set, implementation is capable of data scrambling/descrambling with LFSR.



Device: 2 Function: 1, 5 Offset: 68h Access as a Dword			
Bit	Type	Reset Value	Description
21:20	RO	-	RAS_CAPABILITY. Any of these bits set indicates Alternate Clock RAS capability available and that corresponding control bits in QPI_*_PH_CTR are implemented.
17:16	RO	-	DETERMINISM_SUPPORT. Determinism supported mode of operations. Bit17: If set, Master mode of operation supported. Component Specification or equivalent document should contain the information about PhyLOSynch. Bit16: If set, Slave mode of operation supported.
10:8	RO	-	LINK_WIDTH_CAPABILITY. Bit8: If set, Full Width capable.
7:5	RO	0	DEBUG_CAPABILITY. Bit7: If set, an implementation is not capable of extracting slave electrical parameter from TS.Loopback and apply during the test. Bit6: If set, an implementation is not capable of running in Compliance slave mode as well as transitioning to Loopback.Pattern from Compliance state. Bit5: If set, an implementation is not capable of doing Loopcount Stal
4	RO	0	RETRAIN_GRANULARITY. If set, implementation is capable of 16UI granularity in retraining duration.
3:0	RO	-	PHY_VERSION. This is the Intel QPI Phy version. 0: Current Intel QPI version 0. Rest are reserved.

2.9.2 QPI_0_PH_CTR QPI_1_PH_CTR

Intel QPI Physical Layer Control Register.

Device: 2 Function: 1, 5 Offset: 6Ch Access as a Dword			
Bit	Type	Reset Value	Description
27	RW	0	LA_LOAD_DISABLE. Disables the loading of the effective values of the Intel QPI CSRs when set.
23	RW	0	ENABLE_PRBS. Enables LFSR pattern during bitlock/training. 1 - use pattern in bitlock/retraining. 0 - use clock pattern for bitlock/retraining.
22	RW	0	ENABLE_SCRAMBLE. Enables data scrambling through LFSR. 1 - data scrambled/descrambled with LFSR 0 - data not scrambled/descrambled.
15:14	RW	2	DETERMINISM_MODE. Sets determinism mode of operation. 00 - Non-deterministic initialization. 01 - Slave mode initialization. 10 - Master mode of initialization - valid only if a component can generate its PhyLOSynch.
13	RW	1	DISABLE_AUTO_COMP. Disables automatic entry into compliance. 0 - path from detect.clkterm to compliance is allowed. 1 - path from detect.clkterm to compliance is disabled.
12	RW	0	INIT_FREEZE. When set, freezes the FSM when initialization aborts.



Device: 2 Function: 1, 5 Offset: 6Ch Access as a Dword			
Bit	Type	Reset Value	Description
11	RW	0	DISABLE_ISI_CHECK. Defeature mode to disable ISI checking during Polling.LaneDeskew state.
10:8	RW	0	INIT_MODE. Initialization mode that determines altered initialization modes.
7	RW	0	LINK_SPEED. Identifies slow speed or at-speed operation for the Intel QPI port. 1 - Force direct operational speed initialization. 0 - Slow speed initialization.
5	RW	1	PHYINITBEGIN. Instructs the port to start initialization.
4	RW	0	SINGLE_STEP. Enables single step mode.
3	RW	0	LAT_FIX_CTL. If set, instructs the remote agent to fix the latency.
2	RW	0	BYPASS_CALIBRATION. Indicates the physical layer to bypass calibration.
1	RW	0	RESET_MODIFIER. Modifies soft reset to default reset when set.
0	RW1S	0	PHY_RESET. Physical Layer Reset.

2.9.3

QPI_0_PH_PIS QPI_1_PH_PIS

Intel QPI Physical Layer Initialization Status Register.

Device: 2 Function: 1, 5 Offset: 80h Access as a Dword			
Bit	Type	Reset Value	Description
29	RO	-	GLOBAL_ERROR. Set upon any error detected on the link during Loopback Pattern.
28	RO	-	TEST_BUSY. Test busy bit indicating that a test is in progress.
27	RW1C	0	STATE_HOLD. State machine hold bit for single step and init freeze modes.
26	RO	-	INIT_SPEED. Current initialization speed. 1 - Operational Speed Initialization. 0 - Slow Speed Initialization.
25	RO	-	PORT_RMT_ACK. Port Remote ACK status.
24	RO	-	PORT_TX_RDY. Port Tx Ready status.
20:16	RO	-	RX_STATE. Current state of the local Rx.
12:8	RO	-	TX_STATE. Current state of the local Tx.
1	RW1C	0	CALIBRATION_DONE. Indicates that calibration has been completed for the Intel QPI link.
0	RW1C	0	LINKUP_IDENTIFIER. Link up identifier for the Intel QPI link. Set to 0 during Default Reset. Set to 1 when initialization completes and link enters L0.



2.9.4 QPI_0_PH_PTV QPI_1_PH_PTV

Intel QPI Physical Layer Initialization Primary Timeout Value Register.

Device: 2 Function: 1, 5 Offset: 94h Access as a Dword			
Bit	Type	Reset Value	Description
19:16	RW	0	POLLING_BITLOCK. Exponential count for Polling Bitlock. Timeout value is $2^{(\text{count in this field})} \times 128$ TSL.
11:8	RW	1	INBAND_RESET. Exponential count for Inband_Reset_Init. Time-out value is $2^{(\text{count in this field})} \times 128$ TSL.
3:0	RW	2	DEBOUNCE. Exponential count for debounce.

2.9.5 QPI_0_PH_LDC QPI_1_PH_LDC

Intel QPI Physical Layer Link Determinism Control Register.

Device: 2 Function: 1, 5 Offset: 9Ch Access as a Dword			
Bit	Type	Reset Value	Description
23:16	RW	0	TARGET_LINK_LATENCY. This field specifies the target link latency value in UI that the remote port needs to adjust to.
11:8	RW	5	DRIFT_BUF_DEPTH. The default pointer separation for the Intel QPI Rx PI FIFO.
3:0	RW	2	DRIFT_ALARM_THRESHOLD. Intel QPI RX PI FIFO alarm threshold.



2.9.6 QPI_0_PH_PRT QPI_1_PH_PRT

Intel QPI Periodic Retraining Timing Register

Device: 2 Function: 1, 5 Offset: A4h Access as a Dword			
Bit	Type	Reset Value	Description
22	RW	0	DURATION_GRANULARITY. 1 indicates agent is using 16 UI granularity 0 indicates agent is using 64 UI granularity.
21:14	RW	-	RETRAIN_PKT_CNT. Retraining packet count.
13:10	RW	-	EXP_RETRAIN_INTERVAL. Exponential count for retraining interval. Interval value is multiplied by $2^{(\text{count in this field})}$. Although these values are specified in exponential form, counting still needs to be accurate to single UI.
7:0	RW	-	RETRAIN_INTERVAL. Periodic retraining interval. A value of 0 indicates periodic retraining is disabled. Retraining must be disabled in Slow Mode. Value to be programmed by firmware. Each count represents 1024 UI (16 TSL)

2.9.7 QPI_0_PH_PMR0 QPI_1_PH_PMR0

Intel QPI Physical Layer Power Management Register.

Device: 2 Function: 1, 5 Offset: D0h Access as a Dword			
Bit	Type	Reset Value	Description
27:26	RW	0	LOs_SLEEP_MIN_REM. Remote agent's minimum LOS time. 00 -> 32 UI 01 -> 48 UI 10 -> 64 UI 11 -> 96 UI
21:16	RW	0	LOs_WAKE_REM. Remote agent's LOS wake time in effect. Value is $(\text{field} + 1) * 16 \text{ UI}$.
11:10	RW	-	LOs_SLEEP_MIN. Minimum time local Tx on a port initiating L0s entry should stay in L0s. 00 -> 32 UI 01 -> 48 UI 10 -> 64 UI 11 -> 96 UI
5:0	RW	-	LOs_WAKE. L0s Wake-up time to be used by remote Tx. This parameter value is derived from field value as $(\text{field} + 1) * 16 \text{ UI}$. Field value of 0 (parameter value of 16) means L0s is not supported.



2.9.8 QPI_0_EP_SR QPI_1_EP_SR

Intel QPI Physical Layer Electrical Parameter Select Register. This register enables the equalization coefficient setting functionality of the QPI_[0,1]_EP_MCTR register when QPI_[0,1]_EP_SR is set to 6.

Device: 2 Function: 1, 5 Offset: E0h Access as a Dword			
Bit	Type	Reset Value	Description
23:16	RW	0	EPARAM_SEL. Select electrical parameter. Set to 6 to enable equalization coefficient setting functionality of QPI_[0,1]_EP_MCTR register.

2.9.9 QPI_0_EP_MCTR QPI_1_EP_MCTR

Intel QPI Electrical Parameter Miscellaneous Control Register. This register holds equalization coefficient parameters.

Device: 2 Function: 1, 5 Offset: F4h Access as a Dword			
Bit	Type	Reset Value	Description
31:8	RW	0	MISC_EPARAM_CTL. Miscellaneous electrical-parameter specific control.
7:3	RW	12	TX_EQUALIZATION. Sets the equalization coefficient of the QPI transmitter based on value obtained from SISTA1 simulations.
2	RW	1	EN. Enables or disables custom TEQ setting. 1 - Enable 0 - Disable
1:0	RW	0	RSVD.



2.10 Intel QPI Miscellaneous Registers

2.10.1 QPI_0_PLL_STATUS QPI_1_PLL_STATUS


This register provides the current and available operating conditions for the Intel QPI PLLs.

Device: 2 Function: 1, 5 Offset: 50h Access as a Dword			
Bit	Type	Reset Value	Description
30:24	RO	-	MAX_CCLK_RATIO . Maximum CCLK (The Intel® QuickPath Interconnect Forwarded Clock for at speed operation) supported on this part (Value * 133Mhz).
22:16	RO	-	MIN_CCLK_RATIO . Minimum CCLK (The Intel® QuickPath Interconnect Forwarded Clock for at speed operation) supported on this part (Value * 133Mhz).
14:8	RO	-	CCLK_RATIO_MASK . Mask that will be applied to the QPI_[0,1]_PLL_RATIO.NEXT_PLL_RATIO field on reset to obtain the current ratio (I.E. mask of 1 will force only even ratios; mask of 3 forces every 4th ratio).
6:0	RO	-	CURRENT_CCLK_RATIO . The current CCLK (The Intel® QuickPath Interconnect Forwarded Clock for at speed operation) (Value * 133Mhz).

2.10.2 QPI_0_PLL_RATIO QPI_1_PLL_RATIO

This register holds the next PLL multiplier. The write to one link will affect the mirror port as well as both Intel QPI links. The reads are link specific.

Device: 2 Function: 1, 5 Offset: 54h Access as a Dword			
Bit	Type	Reset Value	Description
6:0	RW	12	NEXT_PLL_RATIO . The next Intel QPI PLL ratio to be adopted.



2.11 Integrated Memory Controller Control Registers

The registers in section 2.11 apply only to processors supporting registered DIMMs.

2.11.1 MC_CONTROL

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	CHANNELO_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 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.



2.11.2 MC_STATUS

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	CHANNELO_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.

2.11.3 MC_SMI_DIMM_ERROR_STATUS

SMI 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.



Device: 3 Function: 0 Offset: 50h Access as a Dword			
Bit	Type	Reset Value	Description
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

2.11.4 MC_SMI_CNTRL

System Management Interrupt control register.

Device: 3 Function: 0 Offset: 54h Access as a Dword			
Bit	Type	Reset Value	Description
16	RW	0	INTERRUPT_SELECT_NMI. NMI enable. Set to enable NMI signaling. Clear to disable NMI signaling. If both NMI and SMI enable bits are set, then only SMI is sent.
15	RW	0	INTERRUPT_SELECT_SMI. SMI enable. Set to enable SMI signaling. Clear to 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_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.



2.11.5 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.

2.11.6 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	RDLCH0. 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	WRLCH0. 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



2.11.7 MC_MAX_DOD

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 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.

2.11.8 MC_RD_CRDT_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 illegal 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 if High or Critical credits are nonzero.
- $CRDT_RD_HIGH + CRDT_RD_CRIT$ 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 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.

2.11.9 MC_CRDT_WR_THLD

Memory Controller Write Credit Thresholds. 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 $MC_Channel_*_WAQ_PARAMS.ISOCENTRYTHRESHHOLD$ 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.



2.11.10 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. 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. 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.

2.11.11 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
12	RO	0	MEMBIST_INPROGRESS. When this bit is asserted by hardware MemTest/MemInit is in progress.
11	RO	0	MEMBIST_CMPLT. When this bit is asserted by hardware MemTest/MemInit is complete.
10	WO	0	RESET_MEMBIST_STATUS. When this bit is written to a 1, the status field MEMBIST_CMPLT is cleared.
9:8	RW	0	CHNL. 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. 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.
5:4	RW	0	RANK. 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.
3:0	RW	0	BANK. 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.




2.12 TAD - Target Address Decoder Registers

2.12.1 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 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.



2.12.2 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 Channel17. 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 Channel16. 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 Channel15. 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 Channel14. 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 Channel13. 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

2.13 Integrated Memory Controller RAS Registers

2.13.1 MC_SSRCONTROL

Scrubbing control. This register allows the enabling of patrol scrubbing and demand scrubbing.

Device: 3 Function: 2 Offset: 48h Access as a Dword			
Bit	Type	Reset Value	Description
14:7	RW	0	SCRATCHPAD. This field is available as a scratchpad for Scrubbing operations.
6	RW	0	DEMAND_SCRUB_EN. Enable Demand Scrubs.
1:0	RW	0	SSR_MODE. Patrol scrub enable. 00: Disable Patrol Scrub 01: Enable Patrol Scrub 10: RSVD.



2.13.2 MC_SCRUB_CONTROL

Contains the Scrub control parameters and status.

Device: 3 Function: 2 Offset: 4Ch Access as a Dword			
Bit	Type	Reset Value	Description
26	RW	0	SCRUBISSUED. When Set, the scrub address registers contain the last scrub address issued.
25	RW	0	ISSUEONCE. When Set, the patrol scrub engine will issue the address in the scrub address registers only once and stop.
24	RW	0	STARTSCRUB. When Set, the Patrol scrub engine will start from the address in the scrub address registers. Once the scrub is issued this bit is reset.
23:0	RW	0	SCRUBINTERVAL. Defines the interval in DCLKS between patrol scrub requests. The calculation for this register to get a scrub to every line in 24 hours is: $((36400)/(\text{memory capacity}/64))/\text{cycle time of DCLK}$. For 512MB at DDR3-800: $(36400/((2^{29})/64))/1.25 \times 10^{-9} = 3471374 = 0x34F80E$.

2.13.3 MC_RAS_ENABLES

RAS enables register.

Device: 3 Function: 2 Offset: 50h Access as a Dword			
Bit	Type	Reset Value	Description
1	RW	0	LOCKSTEPEN. Lockstep enable. When set, channel 0 and 1 are tied together in lockstep. The channel mapper register must be appropriately programmed as well.
0	RW	0	MIRROREN. Mirror mode enable. The channel mapping must be set up before this bit will have an effect on Integrated Memory Controller operation. This changes the error policy and alternates reads between channels.

2.13.4 MC_RAS_STATUS

RAS status register.

Device: 3 Function: 2 Offset: 54h Access as a Dword			
Bit	Type	Reset Value	Description
0	RW	0	REDUNDANCY_LOSS. One channel of a mirrored pair had an uncorrectable error and redundancy has been lost. This bit is set by hardware and must be cleared by software performing a channel reset to regain mirrored status.



2.13.5 MC_SSRSTATUS

Provides the status of the operation specified in MC_SSRCONTROL.SSR_Mode.

Device: 3 Function: 2 Offset: 60h Access as a Dword			
Bit	Type	Reset Value	Description
1	RO	0	INPROGRESS. Patrol Scrub operation in progress. This bit is set by hardware once scrubbing operation has started. It is cleared once operation is complete or fails.
0	RO	0	CMPLT. Patrol Scrub operation complete. Set by hardware once operation is complete. Bit is cleared by hardware when a new operation is enabled.

2.13.6 MC_COR_ECC_CNT_0 MC_COR_ECC_CNT_1 MC_COR_ECC_CNT_2 MC_COR_ECC_CNT_3 MC_COR_ECC_CNT_4 MC_COR_ECC_CNT_5

Per Dimm counters of correctable ECC errors. The register organization is as follows. For example, if there are three DIMMS on the channel, MC_COR_ECC_CNT_0 contains the error counter information for DIMM 0 and DIMM1 on Channel 0. MC_COR_ECC_CNT_1 contains the error counter information for DIMM2 on Channel 0.

The lower 16-bit of MC_COR_ECC_CNT_0 contains the errors for DIMM0 and the upper 16-bit field contains the errors for DIMM1. The lower 16-bit of MC_COR_ECC_CNT_1 contains the errors for DIMM2. The upper 16 bits of MC_COR_ECC_CNT_1 are not used. The same organization applies to Channel 1 and Channel 2.

MC_COR_ECC_CNT_0 : Channel 0 Dimm 0/1
MC_COR_ECC_CNT_1 : Channel 0 Dimm 2/Rsvd
MC_COR_ECC_CNT_2 : Channel 1 Dimm 0/1
MC_COR_ECC_CNT_3 : Channel 1 Dimm 2/Rsvd
MC_COR_ECC_CNT_4 : Channel 2 Dimm 0/1
MC_COR_ECC_CNT_5 : Channel 2 Dimm 2/Rsvd

If there are one or two DIMMS on the channel, the lower 16-bit field of MC_COR_ECC_CNT_0 contains the errors for DIMM0 on Ranks 0 and 1 on Channel 0. The upper 16-bit field contains information for DIMM0 on Ranks 2 and 3 for a quad rank DIMM. The same organization follows for DIMM1 for MC_COR_ECC_CNT_1.

MC_COR_ECC_CNT_0 : Channel 0 Dimm 0 Ranks 0,1/2,3
MC_COR_ECC_CNT_1 : Channel 0 Dimm 1 Ranks 0,1/2,3
MC_COR_ECC_CNT_2 : Channel 1 Dimm 0 Ranks 0,1/2,3
MC_COR_ECC_CNT_3 : Channel 1 Dimm 1 Ranks 0,1/2,3
MC_COR_ECC_CNT_4 : Channel 2 Dimm 0 Ranks 0,1/2,3
MC_COR_ECC_CNT_5 : Channel 2 Dimm 1 Ranks 0,1/2,3



Device: 3 Function: 2 Offset: 80h, 84h, 88h, 8Ch, 90h, 94h Access as a Dword			
Bit	Type	Reset Value	Description
31	RW	0	DIMM1_ERR_OVERFLOW. Correctable error overflow on DIMM 1/Rsvd.
30:16	RW	0	DIMM1_COR_ERR. Correctable error count from DIMM 1/Rsvd.
15	RW	0	DIMM0_ERR_OVERFLOW. Correctable error overflow on DIMM 0/2.
14:0	RW	0	DIMM0_COR_ERR. Correctable error count from DIMM 0/2.

2.14 Integrated Memory Controller Test Registers

2.14.1 MC_TEST_ERR_RCV1

Memory test error recovery and detection. This is another address to access COR_ECC_CNT register. This is the ecc error information for DIMM 2.

Device: 3 Function: 4 Offset: 60h Access as a Dword			
Bit	Type	Reset Value	Description
15	RW	0	DIMM2_ERR_OVERFLOW. Correctable error overflow on DIMM 2.
14:0	RW	0	DIMM2_COR_ERR. Correctable error count from DIMM 2.

2.14.2 MC_TEST_ERR_RCV0

Memory test error recovery and detection. This is another address to access COR_ECC_CNT register. This is the ecc error information for DIMM 0 and DIMM 1.

Device: 3 Function: 4 Offset: 64h Access as a Dword			
Bit	Type	Reset Value	Description
31	RW	0	DIMM1_ERR_OVERFLOW. Correctable error overflow on DIMM 1.
30:16	RW	0	DIMM1_COR_ERR. Correctable error count from DIMM 1.
15	RW	0	DIMM0_ERR_OVERFLOW. Correctable error overflow on DIMM 0.
14:0	RW	0	DIMM0_COR_ERR. Correctable error count from DIMM 0.



2.14.3 MC_TEST_PH_CTR

Memory test Control Register

Device: 3 Function: 4 Offset: 6Ch Access as a Dword			
Bit	Type	Reset Value	Description
10:8	RW	0	INIT_MODE: Initialization Mode Idle: 000 Loopback: 001 Memtest: 110 Meminit: 111

2.14.4 MC_TEST_PH_PIS

Memory test physical layer initialization status

Device: 3 Function: 4 Offset: 80h Access as a Dword			
Bit	Type	Reset Value	Description
29	RO	-	GLOBAL_ERROR: Indication that an error was detected during a memory test.

2.14.5 MC_TEST_PAT_GCTR

Pattern Generator Control

Device: 3 Function: 4 Offset: A8h Access as a Dword			
Bit	Type	Reset Value	Description
28:24	RW	6	EXP_LOOP_CNT: Sets the length of the test, defined as $2^{(EXP_LOOP_CNT)}$
21	RW	0	ERROR_COUNT_STALL: Masks all detected errors until cleared
20	RW1S	0	STOP_TEST: Force exit from Loopback.Pattern
19	RW	0	DRIVE_DC_ZERO: Drive 0 on lanes with PAT_DCD asserted
13:12	RW	0	PATBUF_WD_SEL: Select word within pattern buffer to be written
10:9	RW	0	PATBUF_SEL: Select which pattern buffer will be written when MC_TEST_PAT_BA is written
5	RW	0	IGN_REM_PARAM: Slave will ignore remote parameters transmitted in Loopback.Marker
4	RW	0	ENABLE_LFSR2: Use scrambled output of Pattern Buffer 2
3	RW	0	ENABLE_LFSR1: Use scrambled output of Pattern Buffer 1
2	RW	1	ENABLE_AUTOINV: Inversion pattern register will rotate automatically once per loop
1	RW	0	STOP_ON_ERROR: Exit Loopback.Pattern upon first detected error
0	RW1S	0	START_TEST: Initiate transition to Loopback.Pattern



2.14.6 MC_TEST_PAT_BA

Memory Test Pattern Generator Buffer

Device: 3 Function: 4 Offset: B0h Access as a Dword			
Bit	Type	Reset Value	Description
31:0	RW	0	DATA: 32-bit window into the indirectly-addressed pattern buffer register space.

2.14.7 MC_TEST_PAT_IS

Memory test pattern inversion selection register

Device: 3 Function: 4 Offset: BCh Access as a Dword			
Bit	Type	Reset Value	Description
7:0	RW	1	LANE_INVERT: Per-lane selection of normal or inverted pattern

2.14.8 MC_TEST_PAT_DCD

Memory test DC drive register

Device: 3 Function: 4 Offset: C0h Access as a Dword			
Bit	Type	Reset Value	Description
7:0	RW	0	LANE_DRIVE_DC: Per-lane selection of DC pattern



2.15 Integrated Memory Controller Channel Control Registers

2.15.1 MC_CHANNEL_0_DIMM_RESET_CMD MC_CHANNEL_1_DIMM_RESET_CMD MC_CHANNEL_2_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.

2.15.2 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.



Device: 4, 5, 6 Function: 0 Offset: 54h Access as a Dword			
Bit	Type	Reset Value	Description
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.

2.15.3 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 2ⁿ count.

Bits [24:22] control the logical to physical rank mapping. The Integrated Memory Controller needs to know the location of different ranks in order to drive the proper chip selects (CS#) and Clock Enable (CKE). Each valid combination results in a different mapping of CS or CKE connections to the logical ranks. The table below summarizes the supported combinations.

3DP[24]	SQRP[23]	QRP[22]	Notes
1	0	0	3 DIMMs Per Channel (6ODT/6CS)
0	1	1	Single Quad Rank (2ODT/4CS)
0	0	1	Quad Rank plus another DIMM (4ODT/8CS)
0	0	0	All other configurations.



Register Description

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 RC0 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.



2.15.4 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



2.15.5 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.



2.15.6 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.

2.15.7 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 MR0, 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	MR0. The values to write to MR0 for A15:A0.



2.15.8 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 RC0 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 RC0 register on RDIMMS. This value will be written whenever address inversion is enabled or disabled by hardware. For this reason bit 0 of the RC0 field (bit 16 of this register) will be controlled by hardware. [19:17] will be driven with the RDIMM register write command for RC0.
15:0	RW	0	MR2. The values to write to MR2 for A15:A0.

2.15.9 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



2.15.10 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 1101: RSVD 1110: RSVD 1111: RSVD



Device: 4, 5, 6 Function: 0 Offset: 80h Access as a Dword			
Bit	Type	Reset Value	Description
18:15	RW	0	tdRdTWrr . 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 1101: RSVD 1110: RSVD 1111: RSVD
14:11	RW	0	tdrRdTWrr . 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 1101: RSVD 1110: RSVD 1111: RSVD
10:7	RW	0	tsrRdTWrr . 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 1101: RSVD 1110: RSVD 1111: RSVD

Register Description



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



2.15.11 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.



2.15.12 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.

2.15.13 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_OPPREF. 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.



2.15.14 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.

2.15.15 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 maintenance counter intervals.
7:5	RW	4	tZQCS . 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 spec.
4:0	RW	0	tZQInit . 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 spec.



2.15.16 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 maintenance counter intervals. Register value is (maintenance counter intervals-1). For example, a setting of 0 will produce one maintenance counter interval.

2.15.17 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 $t_{CAS}-1$.
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 $t_{CAS}-t_{WL}$.
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.



2.15.18 MC_CHANNEL_0_ODT_PARAMS2 MC_CHANNEL_1_ODT_PARAMS2 MC_CHANNEL_2_ODT_PARAMS2

The FORCE_ODT fields are directly mapped to pins. When the force bits are set, the corresponding pin on the interface is always driven high regardless of the cycle that is being generated. 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	RSVD .
6	RW	0	RSVD .
5	RW	0	FORCE_ODT5 . Force ODT pin 5 to always be asserted.
4	RW	0	FORCE_ODT4 . Force ODT pin 4 to always be asserted.
3	RW	0	FORCE_ODT3 . Force ODT pin 3 to always be asserted.
2	RW	0	FORCE_ODT2 . Force ODT pin 2 to always be asserted.
1	RW	0	FORCE_ODT1 . Force ODT pin 1 to always be asserted.
0	RW	0	FORCE_ODT0 . Force ODT pin 0 to always be asserted.

2.15.19 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_RD0 . Bit patterns driven out onto ODT pins when Rank0 is read.



2.15.20 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.

2.15.21 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_WR0. Bit patterns driven out onto ODT pins when Rank0 is written.

2.15.22 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.



2.15.23 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.

2.15.24 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.



2.15.25 MC_CHANNEL_0_MAINTENANCE_OPS MC_CHANNEL_1_MAINTENANCE_OPS MC_CHANNEL_2_MAINTENANCE_OPS

This register enables various maintenance operations such as 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 ZQ, RCOMP in increments of maintenance counter intervals.

2.15.26 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.



2.15.27 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.

2.15.28 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.

2.15.29 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.



2.15.30 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.

2.15.31 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	RSVD.
7:0	RW	0	ADAPTIVETIMEOUTCOUNTER. 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.

2.15.32 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.



2.15.33 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.

2.15.34 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.



2.16 Integrated Memory Controller Channel Address Registers

2.16.1 MC_DOD_CH0_0 MC_DOD_CH0_1 MC_DOD_CH0_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.



2.16.2 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 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.



2.16.3 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 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



2.16.4 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


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 table below 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.

2.17 Integrated Memory Controller Channel Rank Registers

2.17.1 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 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].

2.17.2 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. Defines the offset used in the rank interleave. This is a 2's complement value.
3:0	RW	0	RANK. 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.

2.17.3 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. Defines the offset used in the rank interleave. This is a 2's complement value.
3:0	RW	0	RANK. 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.



2.17.4 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. Defines the offset used in the rank interleave. This is a 2's complement value.
3:0	RW	0	RANK. 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.

2.18 Memory Thermal Control

2.18.1 MC_THERMAL_CONTROL0 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. Selects throttling mode. When in lockstep mode, this field should only be non-zero for Channel0. 0: Throttle disabled 1: Open Loop: Throttle when Virtual Temperature is greater than MC_THROTTLE_OFFSET. 2: Closed Loop: Throttle when MC_CLOSED_LOOP.THROTTLE_NOW is set. 3: Closed Loop: Throttle when MC_DDR_THERM_COMMAND.THROTTLE is set and the MC_DDR_THERM pin is asserted OR OLT will be implemented (Condition 1).



2.18.2 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.

2.18.3 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.

2.18.4 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.



2.18.5 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.

2.18.6 MC_COOLING_COEF0 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.



2.18.7 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. When in Lockstep, this field may not be changed when throttling is possible. This includes THROTTLE_NOW or DDR_THERM# pin assertion, depending on throttling mode selected.
4	RW	0	REF_2X_NOW. Direct control of dynamic 2X refresh if MC_THERMAL_CONTROL.THROTTLE_MODE = 2. This bit can be set only when MC_CHANNEL_X_REFRESH_THROTTLE_SUPPORT.ASR_PRESENT bit is set.
3:0	RW	0	THROTTLE_NOW. Throttler Vector to directly control throttling if MC_THERMAL_CONTROL.THROTTLE_MODE = 2.

2.18.8 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.



2.18.9 MC_RANK_VIRTUAL_TEMP0 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 0x40, the value read from the corresponding MC_RANK_VIRTUAL_TEMP register is 0x20.

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.

2.18.10 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 *Intel® Xeon® Processor 5500 Series Datasheet, Volume 1* (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	RSVD.
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	RW1S	0	LOCK. When set, all bits in this register are RO and cannot be written. Reset will clear the lock.



2.18.11 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 *Intel® Xeon® Processor 5500 Series Datasheet, Volume 1* (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

2.19 Integrated Memory Controller Miscellaneous Registers

2.19.1 MC_DIMM_CLK_RATIO_STATUS

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 - 800Mhz 01000 - 1066Mhz 01010 - 1333Mhz
4:0	RO	0	QCLK_RATIO. Current ratio of Qclk. Value - Qclk. 00000 - RSVD 00110 - 800Mhz 01000 - 1066Mhz 01010 - 1333Mhz



2.19.2 MC_DIMM_CLK_RATIO

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 - 800Mhz 01000 - 1066Mhz 01010 - 1333Mhz

S



3 DIMM Population Requirements

3.1 General Population Requirements

The Intel® 5500 platform offers a wide variety of DIMM configurations. Key parameters used in defining various DIMM configurations are listed in [Table 3-1](#).

Table 3-1. Key Parameters for DIMM Configurations

Parameter	Possible Values
# of Channels	1, 2, or 3
# of DIMM Slots per channel	Two DIMM slots or Three DIMM slots
# of DIMMs Populated per channel	1DPC, 2DPC, or 3DPC (required three DIMM slots per channel)
DIMM Type	RDIMM (w/ECC), UDIMM (w/ or w/o ECC) MetaSDRAM* R-DIMM (8 GB module only)
DIMM Raw Cards	RDIMM Raw Cards as defined by JEDEC: A(1Rx8), B (2Rx8), C (1Rx4), D (2Rx4), E/J (2Rx4), F (4Rx4), or H (4Rx8) UDIMM Raw Cards as defined by JEDEC: A (1Rx8), B (2Rx8), C (1Rx16 ¹), D (1Rx8 w/ECC), E (2Rx8 w/ECC)
DIMM Frequencies	DDR3-800, DDR3-1066, or DDR3-1333

Notes:

1. UDIMM Raw Card C(1Rx16) is not supported in RDIMM/UDIMM combo designs (a combo platform can support either RDIMM only or UDIMM only but not a mix of both types).

Following are generic population requirements:

- All DIMMs must be DDR3 DIMMs.
- The Intel® Xeon® processor 5500 series does not support low voltage (1.35V) DDR3 memory. If 1.35V (DDR3L) and 1.50V (DDR3) DIMMs are mixed, the DIMMs will run at 1.50V.
- Registered DIMMs must be ECC only, Unbuffered DIMMs can be ECC or non-ECC.
- Mixing of Registered and Unbuffered DIMMs is not allowed.
- Mixing of MetaSDRAM* R-DIMM with any other DIMM type is not allowed.
- It is allowed to mix ECC and non-ECC Unbuffered DIMMs. The presence of a single non-ECC Unbuffered DIMM will result in disabling ECC functionality.
- DIMMs with different timing parameters can be installed on different slots within the same channel, but only timings that support the slowest DIMM will be applied to all. As a consequence, faster DIMMs will be operated at timings supported by the slowest DIMM populated. The same interface frequency (DDR3-800, DDR3-1066, or DDR3-1333) will be applied to all DIMMs on all channels on the platform (both processors).
- DIMMs with DDR3-1333 speed are allowed only when one DIMM Per Channel (1DPC) is populated. If two 1333 MT/s capable UDIMMs or RDIMMs are detected in the same channel, BIOS would flag this as a warning and force the speed to 1066 MT/s.
- DIMMs with DDR3-1066 speed are allowed only when two DIMMs Per Channel (2DPC) are populated. If three 1066 MT/s capable UDIMMs or RDIMMs are detected in the same channel, BIOS will force the speed to 800 MT/s.



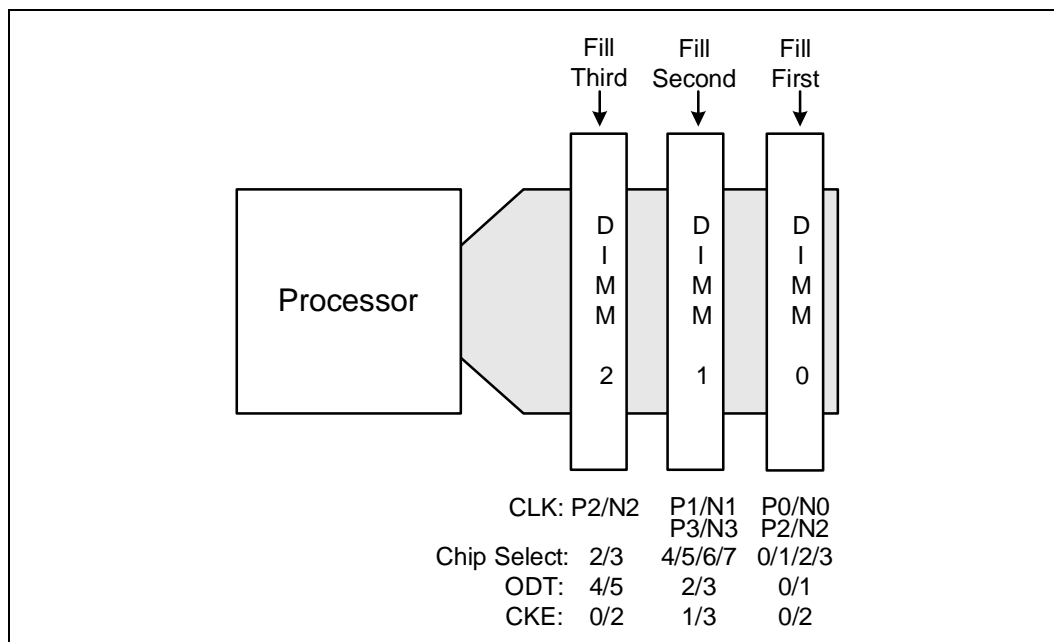
- When one quad rank DIMMs is used, it must be populated in DIMM slot0 (farthest away from the CPU) of a given channel
- Mixing of quad ranks DIMMs (RDIMM Raw Cards F and H) in one channel and three DIMMs in other channel (3DPC) on the same CPU socket is not allowed. If such configuration is detected on a CPU socket, BIOS would flag this as a warning and disable the QR DIMM channel(s).

3.2 Populating DIMMs Within a Channel

3.2.1 DIMM Population for Three Slots per Channel

For three slot per channel configurations, the Intel 5500 platform requires DIMMs within a channel to be populated starting with the DIMMs farthest from the processor in a “fill-farthest” approach (see [Figure 3-1](#)). In addition, when populating a Quad-rank DIMM with a Single- or Dual-rank DIMM in the same channel, the Quad-rank DIMM must be populated farthest from the processor. Note that Quad-rank DIMMs and UDIMMs are not allowed in three slots populated configurations. Intel recommends checking for correct DIMM placement during BIOS initialization. Additionally, Intel strongly recommends that all designs follow the DIMM ordering, command clock, and control signal routing documented in [Figure 3-1](#). This addressing must be maintained to be compliant with the reference BIOS code supplied by Intel. All allowed DIMM population configurations for three slots per channel are shown in [Table 3-2](#) and [Table 3-3](#).

Figure 3-1. DIMM Population within a Channel for Three Slots per Channel



Note: ODT[5:4] is muxed with CS[7:6]#.



Table 3-2. RDIMM Population Configurations within a Channel for Three Slots per Channel

Configuration Number	Maximum Supported Speed ¹	1N or 2N	DIMM2	DIMM1	DIMM0
1	DDR3-1333	1N	Empty	Empty	Single-rank
2	DDR3-1333	1N	Empty	Empty	Dual-rank
3	DDR3-1066	1N	Empty	Empty	Quad-rank
4	DDR3-1066	1N	Empty	Single-rank	Single-rank
5	DDR3-1066	1N	Empty	Single-rank	Dual-rank
6	DDR3-1066	1N	Empty	Dual-rank	Single-rank
7	DDR3-1066	1N	Empty	Dual-rank	Dual-rank
8	DDR3-800	1N	Empty	Single-rank	Quad-rank
9	DDR3-800	1N	Empty	Dual-rank	Quad-rank
10	DDR3-800	1N	Empty	Quad-rank	Quad-rank
11	DDR3-800	1N	Single-rank	Single-rank	Single-rank
12	DDR3-800	1N	Single-rank	Single-rank	Dual-rank
13	DDR3-800	1N	Single-rank	Dual-rank	Single-rank
14	DDR3-800	1N	Dual-rank	Single-rank	Single-rank
15	DDR3-800	1N	Single-rank	Dual-rank	Dual-rank
16	DDR3-800	1N	Dual-rank	Single-rank	Dual-rank
17	DDR3-800	1N	Dual-rank	Dual-rank	Single-rank
18	DDR3-800	1N	Dual-rank	Dual-rank	Dual-rank

Notes:

1. If a DIMM faster than the maximum supported speed is populated, BIOS will force the memory to run at the maximum supported speed.

Table 3-3. UDIMM Population Configurations within a Channel for Three Slots per Channel

Configuration Number	Maximum Supported Speed ¹	1N or 2N	DIMM2	DIMM1	DIMM0
1	DDR3-1333	1N	Empty	Empty	Single-rank
2	DDR3-1333	1N	Empty	Empty	Dual-rank
3	DDR3-1066	2N	Empty	Single-rank	Single-rank
4	DDR3-1066	2N	Empty	Single-rank	Dual-rank
5	DDR3-1066	2N	Empty	Dual-rank	Single-rank
6	DDR3-1066	2N	Empty	Dual-rank	Dual-rank

Notes:

1. If a DIMM faster than the maximum supported speed is populated, BIOS will force the memory to run at the maximum supported speed.



Table 3-4. MetaSDRAM* R-DIMM¹ Population Configurations within a Channel for Three Slots per Channel

Configuration Number	Maximum Supported Speed ²	1N or 2N	DIMM2	DIMM1	DIMM0
1	DDR3-1066	1N	Empty	Empty	Dual-rank
2	DDR3-1066	1N	Empty	Dual-rank	Dual-rank
3	DDR3-1066	1N	Dual-rank	Dual-rank	Dual-rank

Notes:

1. 8 GB DDR3 MetaSDRAM R-DIMM only. Designers considering the support of MetaSDRAM R-DIMM are recommended to review the platform VR design guidelines as the DC/AC load requirement may be different from that of RDIMM/UDIMM.
2. If a DIMM faster than the maximum supported speed is populated, BIOS will force the memory to run at the maximum supported speed.

3.2.2 DIMM Population for Two Slots per Channel

For two slot per channel configurations, the Intel 5500 platform requires DIMMs within a channel to be populated starting with the DIMMs farthest from the processor in a “fill-farthest” approach (see [Figure 3-2](#)). In addition, when populating a Quad-rank DIMM with a Single- or Dual-rank DIMM in the same channel, the Quad-rank DIMM must be populated farthest from the processor. Intel recommends checking for correct DIMM placement during BIOS initialization. Additionally, Intel strongly recommends that all designs follow the DIMM ordering, command clock, and control signal routing documented in [Figure 3-2](#). This addressing must be maintained to be compliant with the reference BIOS code supplied by Intel. All allowed DIMM population configurations for two slots per channel are shown in [Table 3-5](#) and [Table 3-6](#).

Figure 3-2. DIMM Population Within a Channel for Two Slots per Channel

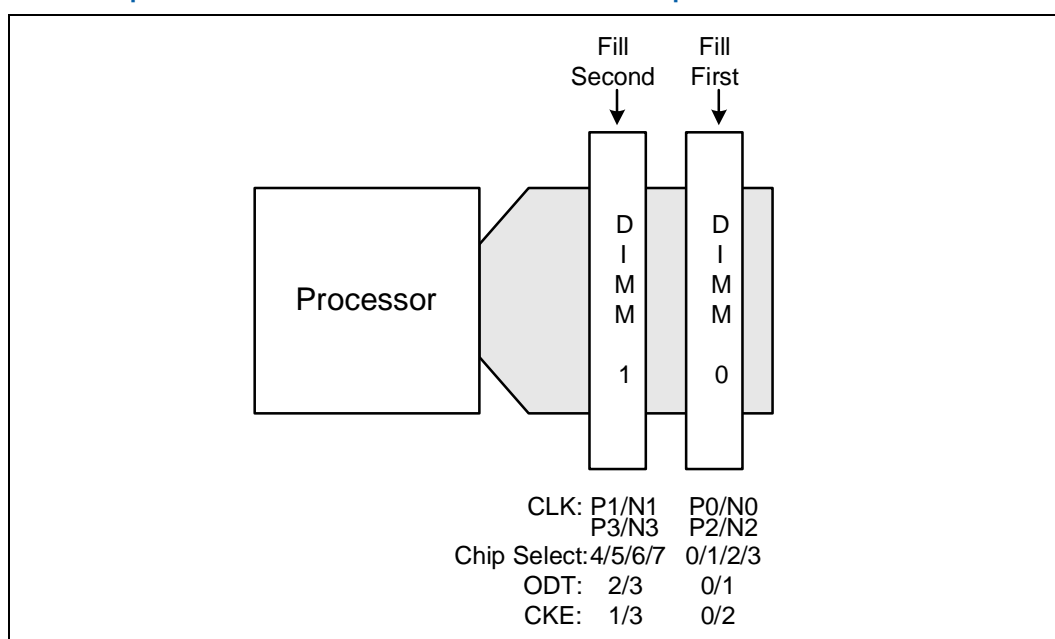




Table 3-5. RDIMM Population Configurations Within a Channel for Two Slots per Channel

Configuration Number	Maximum Supported Speed ¹	1N or 2N	DIMM1	DIMM0
1	DDR3-1333	1N	Empty	Single-rank
2	DDR3-1333	1N	Empty	Dual-rank
3	DDR3-1066	1N	Empty	Quad-rank
4	DDR3-1066	1N	Single-rank	Single-rank
5	DDR3-1066	1N	Single-rank	Dual-rank
6	DDR3-1066	1N	Dual-rank	Single-rank
7	DDR3-1066	1N	Dual-rank	Dual-rank
8	DDR3-800	1N	Single-rank	Quad-rank
9	DDR3-800	1N	Dual-rank	Quad-rank
10	DDR3-800	1N	Quad-rank	Quad-rank

Notes:

1. If a DIMM faster than the maximum supported speed is populated, BIOS will force the memory to run at the maximum supported speed.

Table 3-6. UDIMM Population Configurations within a Channel for Two Slots per Channel

Configuration Number	Maximum Supported Speed ¹	1N or 2N	DIMM1	DIMM0
1	DDR3-1333	1N	Empty	Single-rank
2	DDR3-1333	1N	Empty	Dual-rank
3	DDR3-1066	2N	Single-rank	Single-rank
4	DDR3-1066	2N	Single-rank	Dual-rank
5	DDR3-1066	2N	Dual-rank	Single-rank
6	DDR3-1066	2N	Dual-rank	Dual-rank

Notes:

1. If a DIMM faster than the maximum supported speed is populated, BIOS will force the memory to run at the maximum supported speed.

Table 3-7. MetaSDRAM R-DIMM¹ Population Configurations within a Channel for Two Slots per Channel

Configuration Number	Maximum Supported Speed ²	1N or 2N	DIMM1	DIMM0
1	DDR3-1066	1N	Empty	Dual-rank
2	DDR3-1066	1N	Dual-rank	Dual-rank

Notes:

1. 8 GB DDR3 MetaSDRAM R-DIMM only.
2. If a DIMM faster than the maximum supported speed is populated, BIOS will force the memory to run at the maximum supported speed.



DIMM Population Requirements