Understanding Embedded - Microprocessors

Embedded microprocessors are specialized computing chips designed to perform specific tasks within an embedded system. Unlike general-purpose microprocessors found in personal computers, embedded microprocessors are tailored for dedicated functions within larger systems, offering optimized performance, efficiency, and reliability. These microprocessors are integral to the operation of countless electronic devices, providing the computational power necessary for controlling processes, handling data, and managing communications.

Applications of Embedded - Microprocessors

Embedded microprocessors are utilized across a broad spectrum of applications, making them indispensable in

Details	
Product Status	Active
Core Processor	PowerPC G4
Number of Cores/Bus Width	1 Core, 32-Bit
Speed	1.0GHz
Co-Processors/DSP	Multimedia; SIMD
RAM Controllers	-
Graphics Acceleration	No
Display & Interface Controllers	-
Ethernet	-
SATA	-
USB	-
Voltage - I/O	1.5V, 1.8V, 2.5V
Operating Temperature	0°C ~ 105°C (TA)
Security Features	-
Package / Case	483-BCBGA, FCCBGA
Supplier Device Package	483-FCCBGA (29x29)
Purchase URL	https://www.e-xfl.com/pro/item?MUrl=&PartUrl=mc7457vg1000lc

1 Overview

The MPC7457 is the fourth implementation of the fourth generation (G4) microprocessors from Freescale. The MPC7457 implements the full PowerPC 32-bit architecture and is targeted at networking and computing systems applications. The MPC7457 consists of a processor core, a 512-Kbyte L2, and an internal L3 tag and controller that support a glueless backside L3 cache through a dedicated high-bandwidth interface. The MPC7447 is identical to the MPC7457 except that it does not support the L3 cache interface.

Figure 1 shows a block diagram of the MPC7457. The core is a high-performance superscalar design supporting a double-precision floating-point unit and a SIMD multimedia unit.

The memory storage subsystem supports the MPX bus protocol and a subset of the 60x bus protocol to main memory and other system resources. The L3 interface supports 1, 2, or 4 Mbytes of external SRAM for L3 cache and/or private memory data. For systems implementing 4 Mbytes of SRAM, a maximum of 2 Mbytes may be used as cache; the remaining 2 Mbytes must be private memory.

Note that the MPC7457 is a footprint-compatible, drop-in replacement in a MPC7455 application if the core power supply is 1.3 V.

2 Features

This section summarizes features of the MPC7457 implementation of the PowerPC architecture.

Major features of the MPC7457 are as follows:

- High-performance, superscalar microprocessor
  - As many as four instructions can be fetched from the instruction cache at a time.
  - As many as three instructions can be dispatched to the issue queues at a time.
  - As many as 12 instructions can be in the instruction queue (IQ).
  - As many as 16 instructions can be at some stage of execution simultaneously.
  - Single-cycle execution for most instructions
  - One instruction per clock cycle throughput for most instructions
  - Seven-stage pipeline control
- Eleven independent execution units and three register files
  - Branch processing unit (BPU) features static and dynamic branch prediction
    - 128-entry (32-set, four-way set associative) branch target instruction cache (BTIC), a cache of branch instructions that have been encountered in branch/loop code sequences. If a target instruction is in the BTIC, it is fetched into the instruction queue a cycle sooner than it can be made available from the instruction cache. Typically, a fetch that hits the BTIC provides the first four instructions in the target stream.
    - 2048-entry branch history (BHT) with 2 bits per entry for 4 levels of prediction—not-taken, strongly not-taken, taken, and strongly taken
    - Up to three outstanding speculative branches

- Branch instructions that do not update the count register (CTR) or link register (LR) are often removed from the instruction stream.
- Eight-entry link register stack to predict the target address of Branch Conditional to Link Register (**bclr**) instructions

  - Space must be available in the CQ for an instruction to dispatch (this includes instructions that are assigned a space in the CQ but not in an issue queue)
- Rename buffers
  - 16 GPR rename buffers
  - 16 FPR rename buffers
  - 16 VR rename buffers
- Dispatch unit
  - Decode/dispatch stage fully decodes each instruction
- Completion unit
  - The completion unit retires an instruction from the 16-entry completion queue (CQ) when all instructions ahead of it have been completed, the instruction has finished execution, and no exceptions are pending.
  - Guarantees sequential programming model (precise exception model)
  - Monitors all dispatched instructions and retires them in order
  - Tracks unresolved branches and flushes instructions after a mispredicted branch
  - Retires as many as three instructions per clock cycle
- Separate on-chip L1 instruction and data caches (Harvard architecture)
  - 32-Kbyte, eight-way set associative instruction and data caches
  - Pseudo least recently used (PLRU) replacement algorithm
  - 32-byte (eight-word) L1 cache block
  - Physically indexed/physical tags
  - Cache write-back or write-through operation programmable on a per-page or per-block basis
  - Instruction cache can provide four instructions per clock cycle; data cache can provide four words per clock cycle
  - Caches can be disabled in software.
  - Caches can be locked in software.
  - MESI data cache coherency maintained in hardware
  - Separate copy of data cache tags for efficient snooping
  - L1 cache supports parity generation and checking
  - No snooping of instruction cache except for **icbi** instruction
  - Data cache supports AltiVec LRU and transient instructions
  - Critical double- and/or quad-word forwarding is performed as needed. Critical quad-word forwarding is used for AltiVec loads and instruction fetches. Other accesses use critical double-word forwarding.
- Level 2 (L2) cache interface
  - On-chip, 512-Kbyte, eight-way set associative unified instruction and data cache
  - Fully pipelined to provide 32 bytes per clock cycle to the L1 caches
  - A total nine-cycle load latency for an L1 data cache miss that hits in L2

  - Store merging for multiple store misses to the same line. Only coherency action taken (address-only) for store misses merged to all 32 bytes of a cache block (no data tenure needed).
  - Three-entry finished store queue and five-entry completed store queue between the LSU and the L1 data cache
  - Separate additional queues for efficient buffering of outbound data (such as castouts and write-through stores) from the L1 data cache and L2 cache
- Multiprocessing support features include the following:
  - Hardware-enforced, MESI cache coherency protocols for data cache
  - Load/store with reservation instruction pair for atomic memory references, semaphores, and other multiprocessor operations
- Power and thermal management
  - 1.3-V processor core
  - The following three power-saving modes are available to the system:
    - Nap—Instruction fetching is halted. Only those clocks for the time base, decrementer, and JTAG logic remain running. The part goes into the doze state to snoop memory operations on the bus and back to nap using a $\overline{\text{QREQ}}$/$\overline{\text{QACK}}$ processor-system handshake protocol.
    - Sleep—Power consumption is further reduced by disabling bus snooping, leaving only the PLL in a locked and running state. All internal functional units are disabled.
    - Deep sleep—When the part is in the sleep state, the system can disable the PLL. The system can then disable the SYSCLK source for greater system power savings. Power-on reset procedures for restarting and relocking the PLL must be followed on exiting the deep sleep state.
  - Thermal management facility provides software-controllable thermal management. Thermal management is performed through the use of three supervisor-level registers and an MPC7457-specific thermal management exception.
  - Instruction cache throttling provides control of instruction fetching to limit power consumption
- Performance monitor can be used to help debug system designs and improve software efficiency
- In-system testability and debugging features through JTAG boundary-scan capability
- Testability
  - LSSD scan design
  - IEEE 1149.1 JTAG interface
  - Array built-in self test (ABIST)—factory test only
- Reliability and serviceability
  - Parity checking on system bus and L3 cache bus
  - Parity checking on the L2 and L3 cache tag arrays

Table 1. Microarchitecture Comparison (continued)

Microarchitectural Specs	MPC7457/MPC7447	MPC7455/MPC7445	MPC7450/MPC7451/ MPC7441
Minimum misprediction penalty	6	6	6
Execution Unit Timings (Latency-Throughput)			
Aligned load (integer, float, vector)	3-1, 4-1, 3-1	3-1, 4-1, 3-1	3-1, 4-1, 3-1
Misaligned load (integer, float, vector)	4-2, 5-2, 4-2	4-2, 5-2, 4-2	4-2, 5-2, 4-2
L1 miss, L2 hit latency	9 data/13 instruction	9 data/13 instruction	9 data/13 instruction
SFX (aDd Sub, Shift, Rot, Cmp, logicals)	1-1	1-1	1-1
Integer multiply (32×8, 32×16, 32×32)	3-1, 3-1, 4-2	3-1, 3-1, 4-2	3-1, 3-1, 4-2
Scalar float	5-1	5-1	5-1
VSFX (vector simple)	1-1	1-1	1-1
VCFX (vector complex)	4-1	4-1	4-1
VFPU (vector float)	4-1	4-1	4-1
VPER (vector permute)	2-1	2-1	2-1
MMUs			
TLBs (instruction and data)	128-entry, 2-way	128-entry, 2-way	128-entry, 2-way
Tablewalk mechanism	Hardware + software	Hardware + software	Hardware + software
Instruction BATs/data BATs	8/8	8/8	4/4
L1 I Cache/D Cache Features			
Size	32K/32K	32K/32K	32K/32K
Associativity	8-way	8-way	8-way
Locking granularity	Way	Way	Way
Parity on I cache	Word	Word	Word
Parity on D cache	Byte	Byte	Byte
Number of D cache misses (load/store)	5/1	5/1	5/1
Data stream touch engines	4 streams	4 streams	4 streams
On-Chip Cache Features			
Cache level	L2	L2	L2
Size/associativity	512-Kbyte/8-way	256-Kbyte/8-way	256-Kbyte/8-way
Access width	256 bits	256 bits	256 bits
Number of 32-byte sectors/line	2	2	2
Parity	Byte	Byte	Byte
Off-Chip Cache Support [1]			

5.2 AC Electrical Characteristics

This section provides the AC electrical characteristics for the MPC7457. After fabrication, functional parts are sorted by maximum processor core frequency as shown in Section 1.5.2.1, "Clock AC Specifications," and tested for conformance to the AC specifications for that frequency. The processor core frequency is determined by the bus (SYSCLK) frequency and the settings of the PLL_CFG[0:4] signals. Parts are sold by maximum processor core frequency; see Section 1.11, "Ordering Information."

5.2.1 Clock AC Specifications

Table 8 provides the clock AC timing specifications as defined in Figure 6 and represents the tested operating frequencies of the devices. The maximum system bus frequency, f_{SYSCLK}, given in Table 8 is considered a practical maximum in a typical single-processor system. The actual maximum SYSCLK frequency for any application of the MPC7457 will be a function of the AC timings of the MPC7457, the AC timings for the system controller, bus loading, printed-circuit board topology, trace lengths, and so forth, and may be less than the value given in Table 8. For information regarding the use of spread spectrum clock generators, see Section 9.1.3, "System Bus Clock (SYSCLK) and Spread Spectrum Sources." PLL configuration and bus-to-core multiplier information is found in Section 9.1.1, "Core Clocks and PLL Configuration."

Table 8. Clock AC Timing Specifications

At recommended operating conditions. See Table 4.

Characteristic	Symbol	Maximum Processor Core Frequency								Unit	Notes
		867 MHz		1000 MHz		1200 MHz		1267 MHz			
		Min	Max	Min	Max	Min	Max	Min	Max		
Processor frequency	f_{core}	600	867	600	1000	600	1200	600	1267	MHz	1
VCO frequency	f_{VCO}	1200	1733	1200	2000	1200	2400	1200	2534	MHz	1
SYSCLK frequency	f_{SYSCLK}	33	167	33	167	33	167	33	167	MHz	1, 2
SYSCLK cycle time	t_{SYSCLK}	6.0	30	6.0	30	6.0	30	6.0	30	ns	2
SYSCLK rise and fall time	t_{KR}, t_{KF}	—	1.0	—	1.0	—	1.0	—	1.0	ns	3
SYSCLK duty cycle measured at $OV_{DD}/2$	t_{KHKL}/t_{SYSCLK}	40	60	40	60	40	60	40	60	%	4
SYSCLK cycle-to-cycle jitter		—	150	—	150	—	150	—	150	ps	5, 6

Table 9. Processor Bus AC Timing Specifications [1] (continued)

At recommended operating conditions. See Table 4.

Parameter	Symbol [2]	All Revisions and Speed Grades		Unit	Notes
		Min	Max		
SYSCLK to ARTRY/SHD0/SHD1 high impedance after precharge	t_{KHARPZ}	—	2	t_{SYSCLK}	3, 5, 6, 7

Notes:

1. All input specifications are measured from the midpoint of the signal in question to the midpoint of the rising edge of the input SYSCLK. All output specifications are measured from the midpoint of the rising edge of SYSCLK to the midpoint of the signal in question. All output timings assume a purely resistive 50-Ω load (see Figure 4). Input and output timings are measured at the pin; time-of-flight delays must be added for trace lengths, vias, and connectors in the system.
2. The symbology used for timing specifications herein follows the pattern of $t_{(signal)(state)(reference)(state)}$ for inputs and $t_{(reference)(state)(signal)(state)}$ for outputs. For example, t_{IVKH} symbolizes the time input signals (I) reach the valid state (V) relative to the SYSCLK reference (K) going to the high (H) state or input setup time. And t_{KHOV} symbolizes the time from SYSCLK(K) going high (H) until outputs (O) are valid (V) or output valid time. Input hold time can be read as the time that the input signal (I) went invalid (X) with respect to the rising clock edge (KH) (note the position of the reference and its state for inputs) and output hold time can be read as the time from the rising edge (KH) until the output went invalid (OX).
3. t_{sysclk} is the period of the external clock (SYSCLK) in ns. The numbers given in the table must be multiplied by the period of SYSCLK to compute the actual time duration (in ns) of the parameter in question.
4. According to the bus protocol, TS is driven only by the currently active bus master. It is asserted low then precharged high before returning to high impedance as shown in Figure 6. The nominal precharge width for TS is 0.5 × t_{SYSCLK}, that is, less than the minimum t_{SYSCLK} period, to ensure that another master asserting TS on the following clock will not contend with the precharge. Output valid and output hold timing is tested for the signal asserted. Output valid time is tested for precharge.The high-impedance behavior is guaranteed by design.
5. Guaranteed by design and not tested.
6. According to the bus protocol, ARTRY can be driven by multiple bus masters through the clock period immediately following AACK. Bus contention is not an issue because any master asserting ARTRY will be driving it low. Any master asserting it low in the first clock following AACK will then go to high impedance for one clock before precharging it high during the second cycle after the assertion of AACK. The nominal precharge width for ARTRY is 1.0 t_{SYSCLK}; that is, it should be high impedance as shown in Figure 6 before the first opportunity for another master to assert ARTRY. Output valid and output hold timing is tested for the signal asserted.The high-impedance behavior is guaranteed by design.
7. According to the MPX bus protocol, SHD0 and SHD1 can be driven by multiple bus masters beginning the cycle of TS. Timing is the same as ARTRY, that is, the signal is high impedance for a fraction of a cycle, then negated for up to an entire cycle (crossing a bus cycle boundary) before being three-stated again. The nominal precharge width for SHD0 and SHD1 is 1.0 t_{SYSCLK}. The edges of the precharge vary depending on the programmed ratio of core to bus (PLL configurations).
8. BMODE[0:1] and BVSEL are mode select inputs and are sampled before and after HRESET negation. These parameters represent the input setup and hold times for each sample. These values are guaranteed by design and not tested. These inputs must remain stable after the second sample. See Figure 5 for sample timing.

Figure 4 provides the AC test load for the MPC7457.

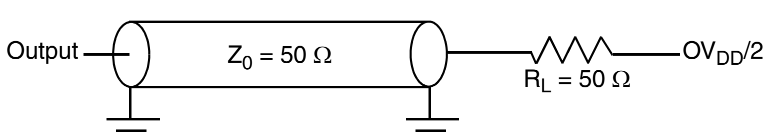

Figure 4. AC Test Load

Figure 5 provides the mode select input timing diagram for the MPC7457.

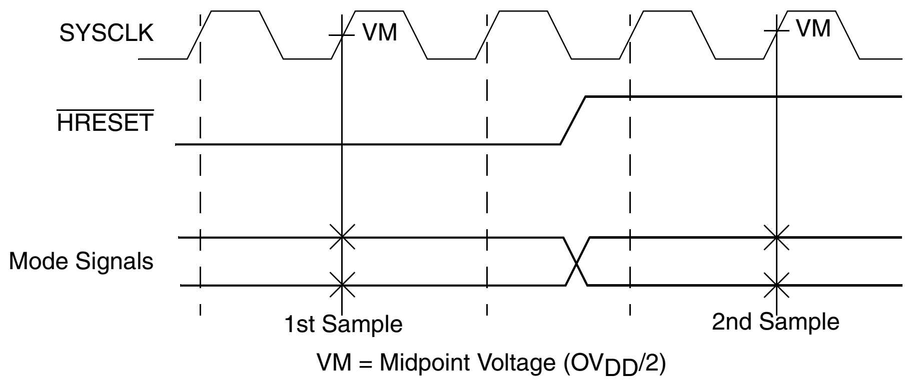

Figure 5. Mode Input Timing Diagram

Figure 6 provides the input/output timing diagram for the MPC7457.

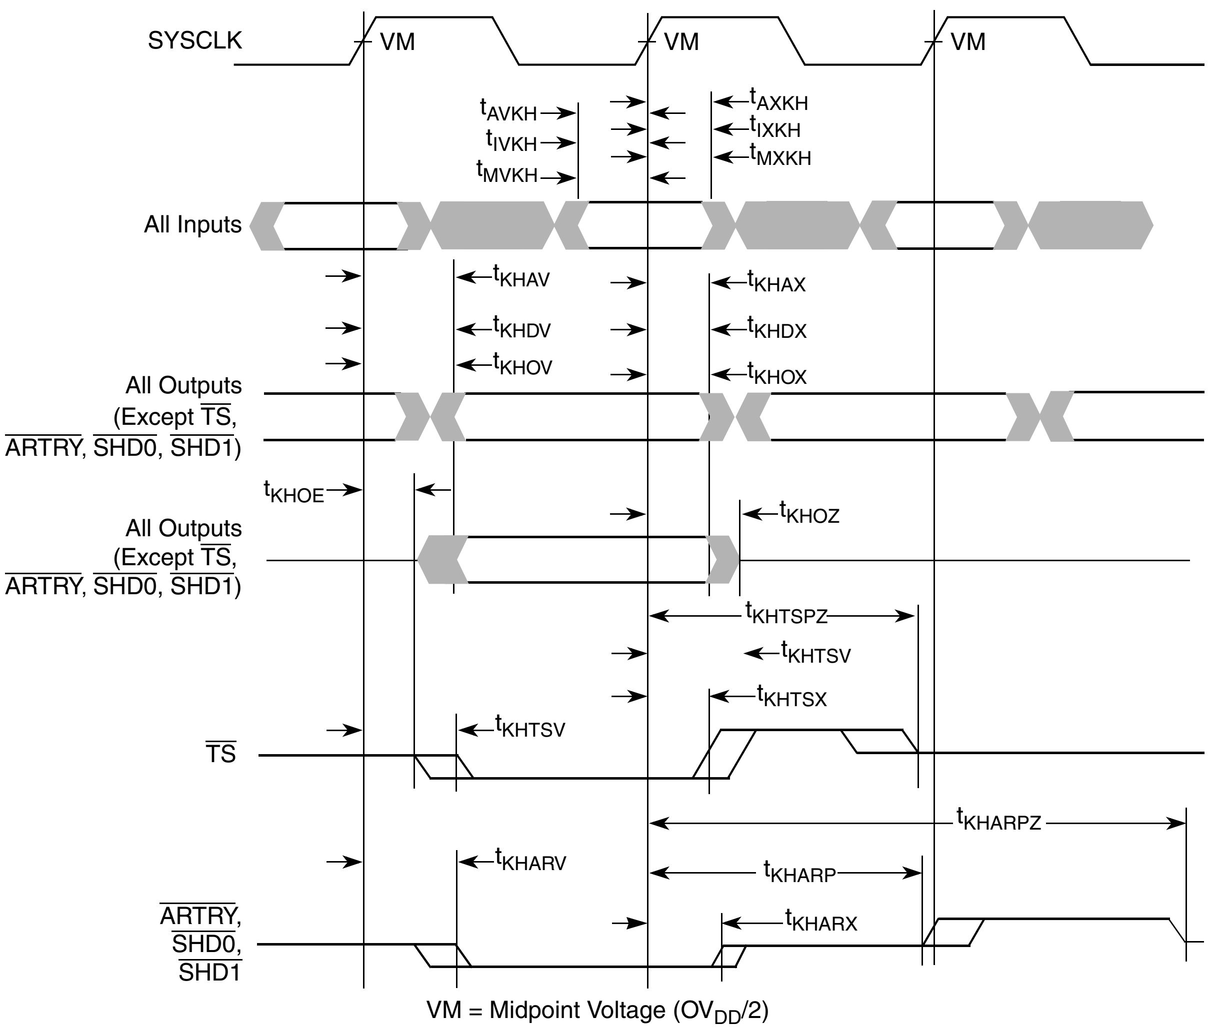

Figure 6. Input/Output Timing Diagram

Table 10. L3_CLK Output AC Timing Specifications (continued)

At recommended operating conditions. See Table 4.

Parameter	Symbol	Device Revision (L3 I/O Voltage) [6]						Unit	Notes
		Rev 1.1. (All I/O Modes) Rev 1.2 (1.5-V I/O Mode)			Rev 1.2 (1.8-, 2.5-V I/O Modes)				
		Min	Typ	Max	Min	Typ	Max		
L3 clock jitter		—	—	± 75	—	—	± 75	ps	5

Notes:

1. The maximum L3 clock frequency (and minimum L3 clock period) will be system dependent. See Section 5.2.3, "L3 Clock AC Specifications," for an explanation that this maximum frequency is not functionally tested at speed by Freescale. The minimum L3 clock frequency and period are f_{SYSCLK} and t_{SYSCLK}, respectively.
2. The nominal duty cycle of the L3 output clocks is 50% measured at midpoint voltage.
3. Maximum possible skew between L3_CLK0 and L3_CLK1. This parameter is critical to the address and control signals which are common to both SRAM chips in the L3.
4. Maximum possible skew between L3_CLK0 and L3_ECHO_CLK1 or between L3_CLK1 and L3_ECHO_CLK3 for PB2 or Late Write SRAM. This parameter is critical to the read data signals because the processor uses the feedback loop to latch data driven from the SRAM, each of which drives data based on L3_CLK0 or L3_CLK1.
5. Guaranteed by design and not tested. The input jitter on SYSCLK affects L3 output clocks and the L3 address, data, and control signals equally and, therefore, is already comprehended in the AC timing and does not have to be considered in the L3 timing analysis. The clock-to-clock jitter shown here is uncertainty in the internal clock period caused by supply voltage noise or thermal effects. This is also comprehended in the AC timing specifications and need not be considered in the L3 timing analysis.
6. L3 I/O voltage mode must be configured by L3VSEL as described in Table 3, and voltage supplied at GV_{DD} must match mode selected as specified in Table 4. See Table 22 for revision level information and part marking.

The L3_CLK timing diagram is shown in Figure 7.

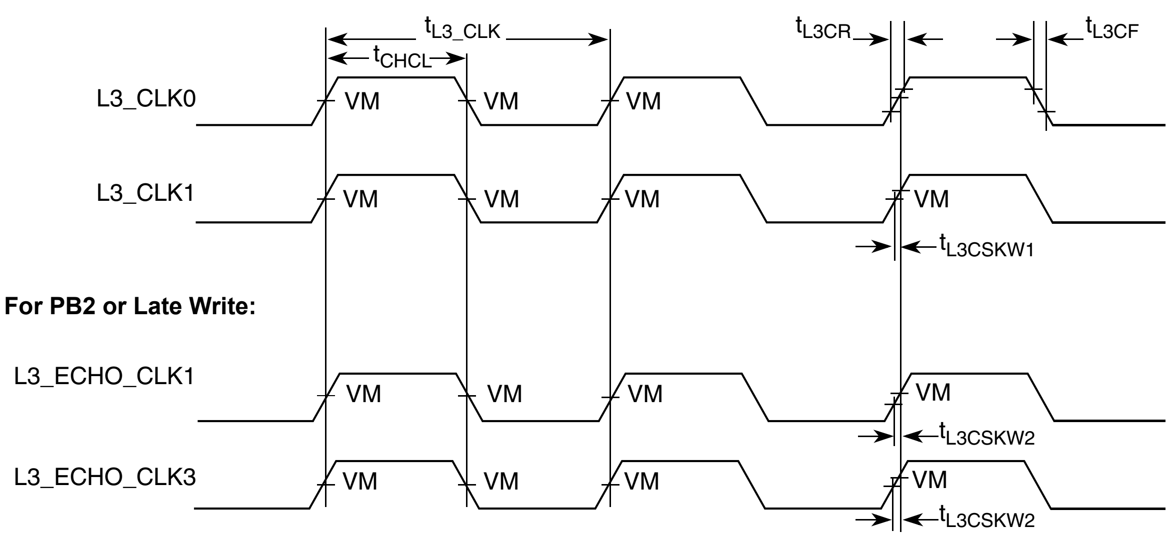

Figure 7. L3_CLK_OUT Output Timing Diagram

Table 12. Effect of L3OHCR Settings on L3 Bus AC Timing (continued)

At recommended operating conditions. See Table 4.

Field Name[1]	Affected Signals	Value	Output Valid Time		Output Hold Time		Unit	Notes
			Parameter Symbol [2]	Change [3]	Parameter Symbol [2]	Change [3]		
L3CLK*n*_OH	All signals latched by SRAM connected to L3_CLK*n*	0b000	t_{L3CHOV}, t_{L3CHDV}, t_{L3CLDV}	0	t_{L3CHOX}, t_{L3CHDX}, t_{L3CLDX}	0	ps	4
		0b001		– 50		– 50		5
		0b010		– 100		– 100		5
		0b011		– 150		– 150		5
		0b100		– 200		– 200		5
		0b101		– 250		– 250		5
		0b110		– 300		– 300		5
		0b111		– 350		– 350		5
L3DOH*n*	L3_DATA[*n*:*n*+7], L3_DP[*n*/8]	0b000	t_{L3CHDV}, t_{L3CLDV}	0	t_{L3CHDX}, t_{L3CLDX}	0	ps	4
		0b001		+ 50		+ 50		
		0b010		+ 100		+ 100		
		0b011		+ 150		+ 150		
		0b100		+ 200		+ 200		
		0b101		+ 250		+ 250		
		0b111		+ 300		+ 300		
		0b111		+ 350		+ 350		

Notes:

1. See the *MPC7450 RISC Microprocessor Family User's Manual* for specific information regarding L3OHCR.
2. See Table 13 and Table 14 for more information.
3. Approximate delay verified by simulation; not tested or characterized.
4. Default value.
5. Increasing values of L3CLK*n*_OH delay the L3_CLK*n* signal, effectively decreasing the output valid and output hold times of all signals latched relative to that clock signal by the SRAM; see Figure 9 and Figure 11.

5.2.4.2 L3 Bus AC Specifications for DDR MSUG2 SRAMs

When using DDR MSUG2 SRAMs at the L3 interface, the parts should be connected as shown in Figure 9. Outputs from the MPC7457 are actually launched on the edges of an internal clock phase-aligned to SYSCLK (adjusted for core and L3 frequency divisors). L3_CLK0 and L3_CLK1 are this internal clock output with 90° phase delay, so outputs are shown synchronous to L3_CLK0 and L3_CLK1. Output valid times are typically negative when referenced to L3_CLK*n* because the data is launched one-quarter period before L3_CLK*n* to provide adequate setup time at the SRAM after the delay-matched address, control, data, and L3_CLK*n* signals have propagated across the printed-wiring board.

Inputs to the MPC7457 are source-synchronous with the CQ clock generated by the DDR MSUG2 SRAMs. These CQ clocks are received on the L3_ECHO_CLK*n* inputs of the MPC7457. An internal circuit delays the incoming L3_ECHO_CLK*n* signal such that it is positioned within the valid data

7 Pinout Listings

Table 16 provides the pinout listing for the MPC7447, 360 CBGA package. Table 17 provides the pinout listing for the MPC7457, 483 CBGA package.

NOTE

This pinout is not compatible with the MPC750, MPC7400, or MPC7410 360 BGA package.

Table 16. Pinout Listing for the MPC7447, 360 CBGA Package

Signal Name	Pin Number	Active	I/O	I/F Select [1]	Notes
A[0:35]	E11, H1, C11, G3, F10, L2, D11, D1, C10, G2, D12, L3, G4, T2, F4, V1, J4, R2, K5, W2, J2, K4, N4, J3, M5, P5, N3, T1, V2, U1, N5, W1, B12, C4, G10, B11	High	I/O	BVSEL	2
$\overline{AACK}$	R1	Low	Input	BVSEL	
AP[0:4]	C1, E3, H6, F5, G7	High	I/O	BVSEL	
$\overline{ARTRY}$	N2	Low	I/O	BVSEL	3
AV_{DD}	A8	—	Input	N/A	
$\overline{BG}$	M1	Low	Input	BVSEL	
$\overline{BMODE0}$	G9	Low	Input	BVSEL	4
$\overline{BMODE1}$	F8	Low	Input	BVSEL	5
$\overline{BR}$	D2	Low	Output	BVSEL	
BVSEL	B7	High	Input	BVSEL	1, 6
$\overline{CI}$	J1	Low	Output	BVSEL	
$\overline{CKSTP\_IN}$	A3	Low	Input	BVSEL	
$\overline{CKSTP\_OUT}$	B1	Low	Output	BVSEL	
CLK_OUT	H2	High	Output	BVSEL	
D[0:63]	R15, W15, T14, V16, W16, T15, U15, P14, V13, W13, T13, P13, U14, W14, R12, T12, W12, V12, N11, N10, R11, U11, W11, T11, R10, N9, P10, U10, R9, W10, U9, V9, W5, U6, T5, U5, W7, R6, P7, V6, P17, R19, V18, R18, V19, T19, U19, W19, U18, W17, W18, T16, T18, T17, W3, V17, U4, U8, U7, R7, P6, R8, W8, T8	High	I/O	BVSEL	
$\overline{DBG}$	M2	Low	Input	BVSEL	
DP[0:7]	T3, W4, T4, W9, M6, V3, N8, W6	High	I/O	BVSEL	
$\overline{DRDY}$	R3	Low	Output	BVSEL	7
DTI[0:3]	G1, K1, P1, N1	High	Input	BVSEL	8
EXT_QUAL	A11	High	Input	BVSEL	9
$\overline{GBL}$	E2	Low	I/O	BVSEL	

Table 16. Pinout Listing for the MPC7447, 360 CBGA Package (continued)

Signal Name	Pin Number	Active	I/O	I/F Select [1]	Notes
$\overline{\text{TEA}}$	L1	Low	Input	BVSEL	
TEST[0:3]	A12, B6, B10, E10	—	Input	BVSEL	12
TEST[4]	D10	—	Input	BVSEL	9
TMS	F1	High	Input	BVSEL	6
$\overline{\text{TRST}}$	A5	Low	Input	BVSEL	6, 14
$\overline{\text{TS}}$	L4	Low	I/O	BVSEL	3
TSIZ[0:2]	G6, F7, E7	High	Output	BVSEL	
TT[0:4]	E5, E6, F6, E9, C5	High	I/O	BVSEL	
$\overline{\text{WT}}$	D3	Low	Output	BVSEL	
V_{DD}	H8, H10, H12, J7, J9, J11, J13, K8, K10, K12, K14, L7, L9, L11, L13, M8, M10, M12	—	—	N/A	

Notes:

1. OV_{DD} supplies power to the processor bus, JTAG, and all control signals; and V_{DD} supplies power to the processor core and the PLL (after filtering to become AV_{DD}). To program the I/O voltage, connect BVSEL to either GND (selects 1.8 V) or to $\overline{\text{HRESET}}$ (selects 2.5 V). If used, the pull-down resistor should be less than 250 Ω. For actual recommended value of V_{in} or supply voltages see Table 4.
2. Unused address pins must be pulled down to GND.
3. These pins require weak pull-up resistors (for example, 4.7 kΩ) to maintain the control signals in the negated state after they have been actively negated and released by the MPC7447 and other bus masters.
4. This signal selects between MPX bus mode (asserted) and 60x bus mode (negated) and will be sampled at $\overline{\text{HRESET}}$ going high.
5. This signal must be negated during reset, by pull up to OV_{DD} or negation by ¬$\overline{\text{HRESET}}$ (inverse of $\overline{\text{HRESET}}$), to ensure proper operation.
6. Internal pull up on die.
7. Ignored in 60x bus mode.
8. These signals must be pulled down to GND if unused, or if the MPC7447 is in 60x bus mode.
9. These input signals are for factory use only and must be pulled down to GND for normal machine operation.
10. This test signal is recommended to be tied to $\overline{\text{HRESET}}$; however, other configurations will not adversely affect performance.
11. These signals are for factory use only and must be left unconnected for normal machine operation.
12. These input signals are for factory use only and must be pulled up to OV_{DD} for normal machine operation.
13. This pin can externally cause a performance monitor event. Counting of the event is enabled via software.
14. This signal must be asserted during reset, by pull down to GND or assertion by $\overline{\text{HRESET}}$, to ensure proper operation.

Table 17. Pinout Listing for the MPC7457, 483 CBGA Package

Signal Name	Pin Number	Active	I/O	I/F Select [1]	Notes
A[0:35]	E10, N4, E8, N5, C8, R2, A7, M2, A6, M1, A10, U2, N2, P8, M8, W4, N6, U6, R5, Y4, P1, P4, R6, M7, N7, AA3, U4, W2, W1, W3, V4, AA1, D10, J4, G10, D9	High	I/O	BVSEL	2
$\overline{\text{AACK}}$	U1	Low	Input	BVSEL	
AP[0:4]	L5, L6, J1, H2, G5	High	I/O	BVSEL	
$\overline{\text{ARTRY}}$	T2	Low	I/O	BVSEL	3

If address or data parity is not used by the system, and the respective parity checking is disabled through HID0, the input receivers for those pins are disabled, and those pins do not require pull-up resistors and should be left unconnected by the system. If all parity generation is disabled through HID0, all parity checking should also be disabled through HID0, and all parity pins may be left unconnected by the system.

The L3 interface does not normally require pull-up resistors. Unused L3_ADDR signals are driven low when the SRAM is configured to be less than 1 M in size via L3CR. For example, L3_ADD[18] will be driven low if the SRAM size is configured to be 2 M; likewise, L3_ADDR[18:17] will be driven low if the SRAM size is configured to be 1 M.

9.7 JTAG Configuration Signals

Boundary-scan testing is enabled through the JTAG interface signals. The $\overline{\text{TRST}}$ signal is optional in the IEEE 1149.1 specification, but is provided on all processors that implement the PowerPC architecture. While it is possible to force the TAP controller to the reset state using only the TCK and TMS signals, more reliable power-on reset performance will be obtained if the $\overline{\text{TRST}}$ signal is asserted during power-on reset. Because the JTAG interface is also used for accessing the common on-chip processor (COP) function, simply tying $\overline{\text{TRST}}$ to $\overline{\text{HRESET}}$ is not practical.

The COP function of these processors allows a remote computer system (typically, a PC with dedicated hardware and debugging software) to access and control the internal operations of the processor. The COP interface connects primarily through the JTAG port of the processor, with some additional status monitoring signals. The COP port requires the ability to independently assert $\overline{\text{HRESET}}$ or $\overline{\text{TRST}}$ in order to fully control the processor. If the target system has independent reset sources, such as voltage monitors, watchdog timers, power supply failures, or push-button switches, the COP reset signals must be merged into these signals with logic.

The arrangement shown in Figure 26 allows the COP port to independently assert $\overline{\text{HRESET}}$ or $\overline{\text{TRST}}$, while ensuring that the target can drive $\overline{\text{HRESET}}$ as well. If the JTAG interface and COP header will not be used, $\overline{\text{TRST}}$ should be tied to $\overline{\text{HRESET}}$ through a 0-Ω isolation resistor so that it is asserted when the system reset signal ($\overline{\text{HRESET}}$) is asserted, ensuring that the JTAG scan chain is initialized during power-on. While Freescale recommends that the COP header be designed into the system as shown in Figure 26, if this is not possible, the isolation resistor will allow future access to $\overline{\text{TRST}}$ in the case where a JTAG interface may need to be wired onto the system in debug situations.

The COP header shown in Figure 26 adds many benefits—breakpoints, watchpoints, register and memory examination/modification, and other standard debugger features are possible through this interface—and can be as inexpensive as an unpopulated footprint for a header to be added when needed.

The COP interface has a standard header for connection to the target system, based on the 0.025" square-post, 0.100" centered header assembly (often called a Berg header). The connector typically has pin 14 removed as a connector key.

There is no standardized way to number the COP header shown in Figure 26; consequently, many different pin numbers have been observed from emulator vendors. Some are numbered top-to-bottom then left-to-right, while others use left-to-right then top-to-bottom, while still others number the pins counter clockwise from pin 1 (as with an IC). Regardless of the numbering, the signal placement recommended in Figure 26 is common to all known emulators.

The $\overline{\text{QACK}}$ signal shown in Figure 26 is usually connected to the PCI bridge chip in a system and is an input to the MPC7457 informing it that it can go into the quiescent state. Under normal operation this occurs during a low-power mode selection. In order for COP to work, the MPC7457 must see this signal asserted (pulled down). While shown on the COP header, not all emulator products drive this signal. If the product does not, a pull-down resistor can be populated to assert this signal. Additionally, some emulator products implement open-drain type outputs and can only drive $\overline{\text{QACK}}$ asserted; for these tools, a pull-up resistor can be implemented to ensure this signal is deasserted when it is not being driven by the tool. Note that the pull-up and pull-down resistors on the $\overline{\text{QACK}}$ signal are mutually exclusive and it is never necessary to populate both in a system. To preserve correct power-down operation, $\overline{\text{QACK}}$ should be merged via logic so that it also can be driven by the PCI bridge.

Tyco Electronics Chip Coolers™ P.O. Box 3668 Harrisburg, PA 17105-3668 Internet: www.chipcoolers.com	800-522-6752
Wakefield Engineering 33 Bridge St. Pelham, NH 03076 Internet: www.wakefield.com	603-635-5102

Ultimately, the final selection of an appropriate heat sink depends on many factors, such as thermal performance at a given air velocity, spatial volume, mass, attachment method, assembly, and cost.

9.8.1 Internal Package Conduction Resistance

For the exposed-die packaging technology, shown in Table 5, the intrinsic conduction thermal resistance paths are as follows:

- The die junction-to-case (actually top-of-die since silicon die is exposed) thermal resistance
- The die junction-to-ball thermal resistance

Figure 28 depicts the primary heat transfer path for a package with an attached heat sink mounted to a printed-circuit board.

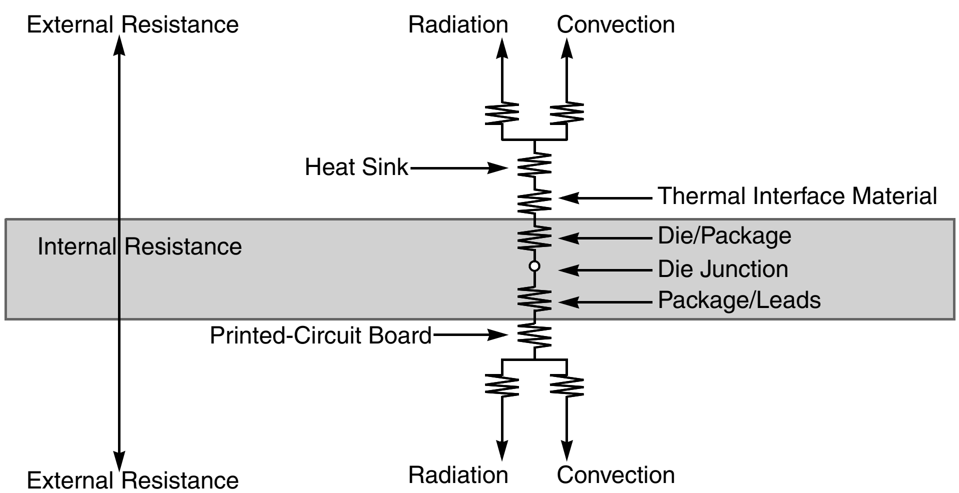

(Note the internal versus external package resistance.)

Figure 28. C4 Package with Heat Sink Mounted to a Printed-Circuit Board

Heat generated on the active side of the chip is conducted through the silicon, through the heat sink attach material (or thermal interface material), and finally to the heat sink where it is removed by forced-air convection.

Because the silicon thermal resistance is quite small, for a first-order analysis, the temperature drop in the silicon may be neglected. Thus, the thermal interface material and the heat sink conduction/convective thermal resistances are the dominant terms.

9.8.2 Thermal Interface Materials

A thermal interface material is recommended at the package lid-to-heat sink interface to minimize the thermal contact resistance. For those applications where the heat sink is attached by spring clip mechanism, Figure 29 shows the thermal performance of three thin-sheet thermal-interface materials (silicone, graphite/oil, floroether oil), a bare joint, and a joint with thermal grease as a function of contact pressure. As shown, the performance of these thermal interface materials improves with increasing contact pressure. The use of thermal grease significantly reduces the interface thermal resistance. That is, the bare joint results in a thermal resistance approximately seven times greater than the thermal grease joint.

Often, heat sinks are attached to the package by means of a spring clip to holes in the printed-circuit board (see Figure 27). Therefore, the synthetic grease offers the best thermal performance, considering the low interface pressure and is recommended due to the high power dissipation of the MPC7457. Of course, the selection of any thermal interface material depends on many factors—thermal performance requirements, manufacturability, service temperature, dielectric properties, cost, etc.

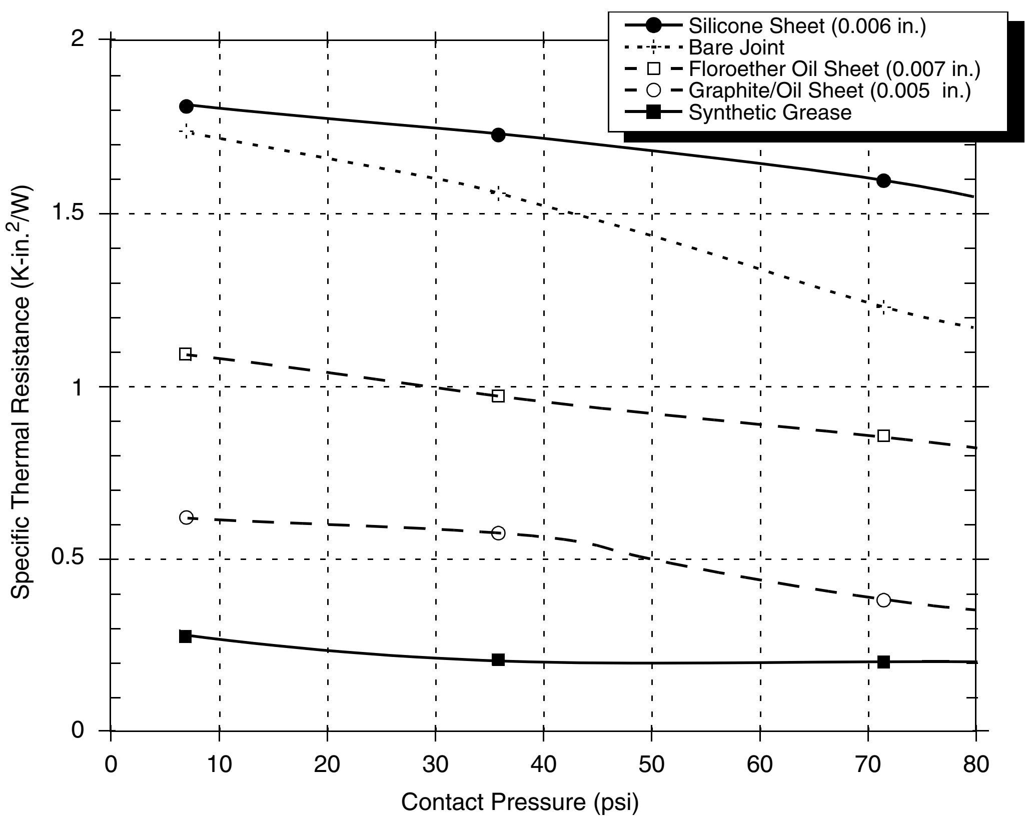

Figure 29. Thermal Performance of Select Thermal Interface Material

The board designer can choose between several types of thermal interface. Heat sink adhesive materials should be selected based on high conductivity, yet adequate mechanical strength to meet equipment shock/vibration requirements. There are several commercially available thermal interfaces and adhesive materials provided by the following vendors:

The Bergquist Company 18930 West 78th St. Chanhassen, MN 55317 Internet: www.bergquistcompany.com	800-347-4572
Chomerics, Inc. 77 Dragon Ct. Woburn, MA 01888-4014 Internet: www.chomerics.com	781-935-4850
Dow-Corning Corporation Dow-Corning Electronic Materials 2200 W. Salzburg Rd. Midland, MI 48686-0997 Internet: www.dow.com	800-248-2481
Shin-Etsu MicroSi, Inc. 10028 S. 51st St. Phoenix, AZ 85044 Internet: www.microsi.com	888-642-7674
Thermagon Inc. 4707 Detroit Ave. Cleveland, OH 44102 Internet: www.thermagon.com	888-246-9050

The following section provides a heat sink selection example using one of the commercially available heat sinks.

9.8.3 Heat Sink Selection Example

For preliminary heat sink sizing, the die-junction temperature can be expressed as follows:

$$T_j = T_I + T_r + (R_{\theta JC} + R_{\theta int} + R_{\theta sa}) \times P_d$$

where:

T_j is the die-junction temperature

T_I is the inlet cabinet ambient temperature

T_r is the air temperature rise within the computer cabinet

$R_{\theta JC}$ is the junction-to-case thermal resistance

$R_{\theta int}$ is the adhesive or interface material thermal resistance

$R_{\theta sa}$ is the heat sink base-to-ambient thermal resistance

P_d is the power dissipated by the device

During operation, the die-junction temperatures (T_j) should be maintained less than the value specified in Table 4. The temperature of air cooling the component greatly depends on the ambient inlet air temperature and the air temperature rise within the electronic cabinet. An electronic cabinet inlet-air temperature (T_a) may range from 30° to 40°C. The air temperature rise within a cabinet (T_r) may be in the range of 5° to 10°C. The thermal resistance of the thermal interface material ($R_{\theta int}$) is typically about 1.5°C/W. For example, assuming a T_a of 30°C, a T_r of 5°C, a CBGA package $R_{\theta JC} = 0.1$, and a typical power consumption (P_d) of 18.7 W, the following expression for T_j is obtained:

Die-junction temperature: $T_j = 30°C + 5°C + (0.1°C/W + 1.5°C/W + \theta_{sa}) \times 18.7$ W

For this example, a $R_{\theta sa}$ value of 2.1°C/W or less is required to maintain the die junction temperature below the maximum value of Table 4.

Though the die junction-to-ambient and the heat sink-to-ambient thermal resistances are a common figure-of-merit used for comparing the thermal performance of various microelectronic packaging technologies, one should exercise caution when only using this metric in determining thermal management because no single parameter can adequately describe three-dimensional heat flow. The final die-junction operating temperature is not only a function of the component-level thermal resistance, but the system-level design and its operating conditions. In addition to the component's power consumption, a number of factors affect the final operating die-junction temperature—airflow, board population (local heat flux of adjacent components), heat sink efficiency, heat sink attach, heat sink placement, next-level interconnect technology, system air temperature rise, altitude, etc.

Due to the complexity and the many variations of system-level boundary conditions for today's microelectronic equipment, the combined effects of the heat transfer mechanisms (radiation, convection, and conduction) may vary widely. For these reasons, we recommend using conjugate heat transfer models for the board, as well as system-level designs.

For system thermal modeling, the MPC7447 and MPC7457 thermal model is shown in Figure 30. Four volumes will be used to represent this device. Two of the volumes, solder ball, and air and substrate, are modeled using the package outline size of the package. The other two, die, and bump and underfill, have the same size as the die. The silicon die should be modeled 9.64 × 11.0 × 0.74 mm with the heat source applied as a uniform source at the bottom of the volume. The bump and underfill layer is modeled as 9.64 × 11.0 × 0.069 mm (or as a collapsed volume) with orthotropic material properties: 0.6 W/(m • K) in the xy-plane and 2 W/(m • K) in the direction of the z-axis. The substrate volume is 25 × 25 × 1.2 mm (MPC7447) or 29 × 29 × 1.2 mm (MPC7457), and this volume has 18 W/(m • K) isotropic conductivity. The solder ball and air layer is modeled with the same horizontal dimensions as the substrate and is 0.9 mm thick. It can also be modeled as a collapsed volume using orthotropic material properties: 0.034 W/(m • K) in the xy-plane direction and 3.8 W/(m • K) in the direction of the z-axis.

Table 22. Part Numbering Nomenclature

MC	***74x7***	***xx***	***nnnn***	**L**	***x***
Product Code	**Part Identifier**	**Package**	**Processor Frequency [1]**	**Application Modifier**	**Revision Level**
PPC [2] MC	7457 7447	RX = CBGA	867 1000 1200 1267	L: 1.3 V ± 50 mV 0° to 105°C	B: 1.1; PVR = 8002 0101
MC	7457	RX = CBGA VG = RoHS BGA	867 1000 1200 1267		C: 1.2; PVR = 8002 0102

Notes:

1. Processor core frequencies supported by parts addressed by this specification only. Parts addressed by a hardware specification addendum may support other maximum core frequencies.
2. The P prefix in a Freescale part number designates a “Pilot Production Prototype” as defined by Freescale SOP 3-13. These parts have only preliminary reliability and characterization data. Before pilot production prototypes may be shipped, written authorization from the customer must be on file in the applicable sales office acknowledging the qualification status and the fact that product changes may still occur while shipping pilot production prototypes.

10.2 Part Numbers Not Fully Addressed by This Document

Parts with application modifiers or revision levels not fully addressed are described in a separate addendum, which supplement and supersede this hardware specification. As such parts are released, these specifications will be listed in this section.

Table 23. Part Numbers Addressed by MPC74x7RX*nnnn*N*x* Series Hardware Specifications Addendum (Document Order No. MPC7457ECS01AD)

MC	***74x7***	***xx***	***nnnn***	**N**	***x***
Product Code	**Part Identifier**	**Package**	**Processor Frequency**	**Application Modifier**	**Revision Level**
PPC	7457	RX = CBGA	1000 867 733 600	N: 1.1 V ± 50 mV 0° to 105°C	B: 1.1; PVR = 8002 0101
	7447		1000 867		
MC	7447		1000 867 733 600		B: 1.1; PVR = 8002 0101
	7457	RX = CBGA VG = RoHS BGA	1000 867 733 600		C: 1.2; PVR = 8002 0102

How to Reach Us:

Home Page:
www.freescale.com
Web Support:
http://www.freescale.com/support

USA/Europe or Locations Not Listed:
Freescale Semiconductor
Technical Information Center, EL516
2100 East Elliot Road
Tempe, Arizona 85284
+1-800-521-6274 or +1-480-768-2130
www.freescale.com/support

Europe, Middle East, and Africa:
Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296 380 456 (English)
+46 8 52200080 (English)
+49 89 92103 559 (German)
+33 1 69 35 48 48 (French)
www.freescale.com/support

Japan:
Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku
Tokyo 153-0064
Japan
0120 191014 or
+81 3 5437 9125
support.japan@freescale.com

Asia/Pacific:
Freescale Semiconductor China Ltd.
Exchange Building 23F
No. 118 Jianguo Road
Chaoyang District
Beijing 100022
China
+86 10 5879 8000
support.asia@freescale.com

Document Number: MPC7457EC
Rev. 8
04/2013