

MPC8540 PowerQUICC III™ Integrated Host Processor Product Brief

The MPC8540 integrates a PowerPC™ processor core with system logic required for networking, storage, and general-purpose embedded applications. The MPC8540 is a member of a growing family of products that combine system-level support for industry standard interfaces to processors that implement the PowerPC architecture. This chapter provides a high-level description of the features and functionality of the MPC8540 integrated microprocessor.

1 Introduction

The MPC8540 uses the e500 core and RapidIO interconnect technology to balance processor performance with I/O system throughput. The e500 core implements the enhanced Book E instruction set architecture and provides unprecedented levels of hardware and software debugging support.

In addition, the MPC8540 offers 256 Kbytes of L2 cache, 2 integrated 10/100/1Gb three-speed Ethernet controllers (TSECs), a 10/100 maintenance port, a DDR SDRAM memory controller, a 64-bit PCI/PCI-X controller, an 8-bit RapidIO port, a programmable interrupt controller, an I²C controller, a 4-channel DMA controller, a general-purpose I/O port, and 2 universal asynchronous receiver/transmitters (DUART). The high level of integration in the MPC8540 simplifies board design and offers significant bandwidth and performance.

2 MPC8540 Overview

The following section provides a high-level overview of the features of the MPC8540.

Figure 1 shows the major functional units in the MPC8540.




Figure 1. MPC8540 Block Diagram

2.1 Key Features

The following is an overview of the MPC8540 feature set.

- High-performance, 32-bit Book E-enhanced core that implements the PowerPC architecture
 - 32-Kbyte L1 instruction cache and 32-Kbyte L1 data cache with parity protection. Caches can be locked entirely or on a per-line basis, with separate locking for instructions and data.
 - Signal-processing engine (SPE) auxiliary processing unit (APU) provides an extensive instruction set for vector (64-bit) integer, single-precision floating-point, and fractional operations. These instructions use both the upper and lower words of the 64-bit GPRs as they are defined by the SPE APU.
 - The single-precision floating-point (SPFP) APU provides an instruction set for single-precision (32-bit) floating-point instructions.
 - Memory management unit (MMU) especially designed for embedded applications
 - Enhanced hardware and software debug support
 - Performance monitor facility (similar to but different from the MPC8540 performance monitor described in the *MPC8540 PowerQUICC III Integrated Host Processor Reference Manual*.)

The e500 defines features that are not implemented on the MPC8540. It also generally defines some features that the MPC8540 implements more specifically. An understanding of these differences can be critical to ensure proper operation. These differences are summarized in Section 5.14, “MPC8540 Implementation Details,” in the reference manual.

Section 3.1, “e500 Core Overview,” in the reference manual includes a comprehensive list of e500 core features.

- 256-Kbyte L2 cache/SRAM
 - Can be configured as follows:
 - Full cache mode (256-Kbyte cache)
 - Full memory-mapped SRAM mode (256-Kbyte SRAM mapped as a single 256-Kbyte block or two 128-Kbyte blocks)
 - Half SRAM and half cache mode (128-Kbyte cache and 128-Kbyte memory-mapped SRAM)
 - Full error checking and correction (ECC) support on 64-bit boundary in both cache and SRAM modes
 - Cache mode supports instruction caching, data caching, or both
 - External masters can force data to be allocated into the cache through programmed memory ranges or special transaction types (stashing).
 - Eight-way set-associative cache organization (1024 sets of 32-byte cache lines)
 - Supports locking the entire cache or selected lines. Individual line locks are set and cleared through Book E instructions or by externally mastered transactions.
 - Global locking and flash clearing done through writes to L2 configuration registers
 - Instruction and data locks can be flash cleared separately
 - Read and write buffering for internal bus accesses
 - SRAM features include the following:
 - I/O devices access SRAM regions by marking transactions as snooperable (global).
 - Regions can reside at any aligned location in the memory map.
 - Byte-accessible ECC is protected using read-modify-write transactions accesses for smaller than cache-line accesses
- Address translation and mapping unit (ATMU)
 - Eight local access windows define mapping within local 32-bit address space
 - Inbound and outbound ATMUs map to larger external address spaces
 - Three inbound windows plus a configuration window on PCI/PCI-X
 - Four inbound windows plus a default and configuration window on RapidIO
 - Four outbound windows plus default translation for PCI
 - Eight outbound windows plus default translation for RapidIO
- DDR memory controller
 - Programmable timing supporting DDR-1 SDRAM
 - 64-bit data interface, up to 333-MHz data rate
 - Four banks of memory supported, each up to 1 Gbyte
 - DRAM chip configurations from 64 Mbits to 1 Gbit with x8/x16 data ports
 - Full ECC support
 - Page mode support (up to 16 simultaneous open pages)
 - Contiguous or discontinuous memory mapping
 - Read-modify-write support for RapidIO atomic increment, decrement, set, and clear transactions
 - Sleep mode support for self-refresh SDRAM

- Supports auto refreshing
- On-the-fly power management using CKE signal
- Registered DIMM support
- Fast memory access through JTAG port
- 2.5-V SSTL2 compatible I/O
- RapidIO interface unit
 - 8-bit RapidIO I/O and messaging protocols
 - Source-synchronous double data rate (DDR) interfaces
 - Supports small type systems (small domain, 8-bit device ID)
 - Supports four priority levels (ordering within a level)
 - Reordering across priority levels
 - Maximum data payload of 256 bytes per packet
 - Packet pacing support at the physical layer
 - CRC protection for packets
 - Supports atomic operations increment, decrement, set, and clear
 - LVDS signaling
- RapidIO-compliant message unit
 - One inbound data message structure (inbox)
 - One outbound data message structure (outbox)
 - Supports chaining and direct modes in the outbox
 - Support of up to 16 packets per message
 - Support of up to 256 bytes per packet and up to 4 Kbytes of data per message
 - Supports one inbound doorbell message structure
- Programmable interrupt controller (PIC)
 - Programming model is compliant with the OpenPIC architecture.
 - Supports 16 programmable interrupt and processor task priority levels
 - Supports 12 discrete external interrupts
 - Supports 4 message interrupts with 32-bit messages
 - Supports connection of an external interrupt controller such as the 8259 programmable interrupt controller
 - Four global high resolution timers/counters that can generate interrupts
 - Supports 22 other internal interrupt sources
 - Supports fully nested interrupt delivery
 - Interrupts can be routed to external pin for external processing.
 - Interrupts can be routed to the e500 core's standard or critical interrupt inputs.
 - Interrupt summary registers allow fast identification of interrupt source.
- I²C controller
 - Two-wire interface
 - Multiple master support
 - Master or slave I²C mode support

- On-chip digital filtering rejects spikes on the bus
- Boot sequencer
 - Optionally loads configuration data from serial ROM at reset through the I²C interface
 - Can be used to initialize configuration registers and/or memory
 - Supports extended I²C addressing mode
 - Data integrity checked with preamble signature and CRC
- DUART
 - Two 4-wire interfaces (SIN, SOUT, $\overline{\text{RTS}}$, $\overline{\text{CTS}}$)
 - Programming model compatible with the original 16450 UART and the PC16550D
- 10/100 fast Ethernet controller (FEC)
 - Operates at 10 to 100 megabits per second (Mbps) as a device debug and maintenance port
- Local bus controller (LBC)
 - Multiplexed 32-bit address and data operating at up to 166 MHz
 - Eight chip selects support eight external slaves
 - Four- and eight-beat burst transfers
 - The 32-, 16-, and 8-bit port sizes are controlled by an on-chip memory controller.
 - Three protocol engines available on a per chip select basis:
 - General-purpose chip select machine (GPCM)
 - Three user programmable machines (UPMs)
 - Dedicated single data rate SDRAM controller
 - Parity support
 - Default boot ROM chip select with configurable bus width (8-,16-, or 32-bit)
- Two three-speed (10/100/1Gb) Ethernet controllers (TSECs)
 - Dual IEEE 802.3, 802.3u, 802.3x, 802.3z, 802.3ac, 802.3ab compliant controllers
 - Support for different Ethernet physical interfaces:
 - 10/100/1Gb IEEE 802.3 GMII
 - 10/100 Mbps IEEE 802.3 MII
 - 10-Mbps IEEE 802.3 MII
 - 1-Gbps IEEE 802.3z TBI
 - 10/100/1Gb RGMII/RTBI
 - Full- and half-duplex support
 - Buffer descriptors are backward compatible with MPC8260 and MPC860T 10/100 programming models
 - 9.6-Kbyte jumbo frame support
 - RMON statistics support
 - 2-Kbyte internal transmit and receive FIFOs
 - MII management interface for control and status
 - Programmable CRC generation and checking
 - Ability to force allocation of header information and buffer descriptors into L2 cache

- OCeaN switch fabric
 - Four-port crossbar packet switch
 - Reorders packets from a source based on priorities
 - Reorders packets to bypass blocked packets
 - Implements starvation avoidance algorithms
 - Supports packets with payloads of up to 256 bytes
- Integrated DMA controller
 - Four-channel controller
 - All channels accessible by both the local and remote masters
 - Extended DMA functions (advanced chaining and striding capability)
 - Support for scatter and gather transfers
 - Misaligned transfer capability
 - Interrupt on completed segment, link, list, and error
 - Supports transfers to or from any local memory or I/O port
 - Selectable hardware-enforced coherency (snoop/no-snoop)
 - Ability to start and flow control each DMA channel from external 3-pin interface
 - Ability to launch DMA from single write transaction
- PCI/PCI-X controller
 - PCI 2.2 and PCI-X 1.0 compatible
 - 64- or 32-bit PCI port supports at 16 to 66 MHz
 - 64-bit PCI-X support up to 133 MHz
 - Host and agent mode support
 - 64-bit dual address cycle (DAC) support
 - PCI-X supports multiple split transactions
 - Supports PCI-to-memory and memory-to-PCI streaming
 - Memory prefetching of PCI read accesses
 - Supports posting of processor-to-PCI and PCI-to-memory writes
 - PCI 3.3-V compatible
 - Selectable hardware-enforced coherency
- Power management
 - Fully static 1.2-V CMOS design with 3.3- and 2.5-V I/O
 - Supports power saving modes: doze, nap, and sleep
 - Employs dynamic power management, which automatically minimizes power consumption of blocks when they are idle
- System performance monitor
 - Supports eight 32-bit counters that count the occurrence of selected events
 - Ability to count up to 512 counter-specific events
 - Supports 64 reference events that can be counted on any of the 8 counters
 - Supports duration and quantity threshold counting

- Burstiness feature that permits counting of burst events with a programmable time between bursts
- Triggering and chaining capability
- Ability to generate an interrupt on overflow
- System access port
 - Uses JTAG interface and a TAP controller to access entire system memory map
 - Supports 32-bit accesses to configuration registers
 - Supports cache-line burst accesses to main memory
 - Supports large block (4-Kbyte) uploads and downloads
 - Supports continuous bit streaming of entire block for fast upload and download
- IEEE 1149.1-compliant, JTAG boundary scan
- 783 FC-PBGA package

3 MPC8540 Architecture Overview

The following sections describe the major functional units of the MPC8540.

3.1 e500 Core Overview

The MPC8540 uses the e500 microprocessor core complex. The e500 core has an internal PLL that allows independent optimization of the operating frequencies. The core frequencies are derived from either the primary PCI clock input or an external oscillator. For information regarding the e500 core refer to the following documents:

- *EREF: A Reference for Freescale Semiconductor Book E and the e500 Core*
- *PowerPC e500 Core Complex Reference Manual*
- *PowerPC e500 Application Binary Interface User's Guide*

NOTE

The e500 defines features that are not implemented on the MPC8540. It also generally defines some features that the MPC8540 implements more specifically. An understanding of these differences can be critical to ensure proper operation. These differences are summarized in Section 5.14, “MPC8540 Implementation Details,” in the *MPC8540 Integrated Processor Reference Manual*.

The following is a brief list of some of the key features of the e500 core complex:

- Implements full Book E 32-bit architecture
- Implements additional instructions, registers, and interrupts defined by APUs. The SPE provides an extensive instruction set for 64-bit vector integer, single-precision floating-point, and fractional operations. The SPFP APU provides scalar (32-bit) single-precision, floating-point instructions.

NOTE

The SPE APU and SPFP APU functionality will be implemented in the MPC8540, the MPC8560 and in their derivatives (that is, in all PowerQUICC III devices). However, these instructions will not be supported in devices subsequent to PowerQUICC III. Freescale Semiconductor strongly recommends that use of these instructions be confined to libraries and device drivers. Customer software that uses SPE or SPFP APU instructions at the assembly level or that uses SPE intrinsics will require rewriting for upward compatibility with next-generation PowerQUICC devices.

Freescale offers a lib_moto_e500 library that uses SPE and SPFP APU instructions. Freescale will also provide future libraries to support next generation PowerQUICC devices.

- L1 cache structure
 - 32-Kbyte, 32-byte line, eight-way set-associative instruction cache
 - 32-Kbyte, 32-byte line, eight-way set-associative data cache
 - 1.5-cycle cache array access, 3-cycle load-to-use latency
 - Pseudo-LRU replacement algorithm
 - Copy-back data cache
- Dual-dispatch superscalar
- Precise exception handling
- Seven-stage pipeline control
- Instruction unit
 - Twelve-entry instruction queue
 - Full hardware detection of interlocks
 - Dispatch up to two instructions per cycle
 - Dispatch serialization control
 - Register dependency resolution and renaming
- Branch unit (BU)
 - Dynamic branch prediction
 - Two-entry branch instruction queue (BIQ)
 - Executes all branch and CR logical instruction
- Completion unit
 - As many as 14 instructions allowed in 14-entry completion queue
 - In-order retirement of up to two instructions per cycle
 - Completion and refetch serialization control
 - Synchronization for all instruction flow changes—interrupts and mispredicted branches
- Two simple execution units that perform the following:
 - Single-cycle add and subtract
 - Single-cycle shift and rotate
 - Single-cycle logical operations
 - Supports integer signal processing operations

- Multiple-cycle execution unit (MU)
 - Four-cycle latency for integer and floating-point multiplication (including integer, fractional, and both vector and scalar floating-point multiply instructions).
 - Variable-latency divide: 4, 11, 19, and 35 cycles for all Book E, SPE, and SPFP divide instructions. Note that the MU allows divide instructions to bypass the second two MU pipeline stages, freeing those stages for other MU instructions to execute in parallel.
 - Four-cycle floating-point multiply
 - Four-cycle floating-point add and subtract
- Signal processing engine APU (SPE APU). The SIMD capability provided by the 64-bit execution units (MIU, LSU, SIU1) is not a separate execution unit. The hardware that executes 32-bit Book E instructions also executes the lower half of 64-bit SPU instructions.
 - Single-cycle integer add and subtract
 - Single-cycle logical operations
 - Single-cycle shift and rotate
 - Four-cycle integer pipelined multiplies
 - 4-, 11-, 19-, and 35-cycle integer divides
 - Four-cycle single instruction multiple data (SIMD) pipelined multiply-accumulate (MAC)
 - 64-bit accumulator for MAC operations
 - Single-precision floating-point operations
- Load/store unit (LSU)
 - Three-cycle load latency
 - Fully pipelined
 - Four-entry load queue allows up to four load misses before stalling
 - Can continue servicing load hits when load queue is full
 - Six-entry store queue allows full pipelining of stores
- Cache coherency
 - Bus support for hardware-enforced coherency (bus snooping)
- Core complex bus (CCB)
 - High-speed, on-chip local bus with data tagging
 - 32-bit address bus
 - 60x-like address protocol with address pipelining and retry/copyback
 - Two general-purpose read data, one write data bus
 - 128-bit data plus parity/tags (each data bus)
 - Supports out-of-order reads, in-order writes
 - Little to no data bus arbitration logic required for native systems
 - Easily adaptable to 60x-like environments
 - Supports one-level pipelining of addresses with address-retry responses
- Extended exception handling
 - Supports Book E interrupt model
 - Interrupt vector prefix register (IVPR)

- Vector offset registers (IVORs) 0–15 as defined in Book E, plus e500-defined IVORs 32–35
- Exception syndrome register (ESR)
- Book E–defined preempting critical interrupt, including critical interrupt status registers (CSRR0 and CSRR1) and an **rfci** instruction
- e500-specific interrupts not defined in Book E architecture
 - SPE APU unavailable exception
 - Floating-point data exception
 - Floating-point round exception
 - Performance monitor
- Memory management unit (MMU)
 - Data L1 MMU
 - Four-entry, fully-associative TLB array for variable-sized pages
 - 64-entry, four-way set-associative TLB for 4-Kbyte pages
 - Instruction L1 MMU
 - Four-entry, fully-associative TLB array for variable-sized pages
 - 64-entry, four-way set-associative TLB for 4-Kbyte pages
 - Unified L2 MMU
 - 16-entry, fully-associative TLB array for variable-sized pages
 - 256-entry, two-way set-associative TLB for 4-Kbyte pages
 - Software reload for TLBs
 - Virtual memory support for as much as 4 Gbytes (2^{32}) of virtual memory
 - Real memory support for as much as 4 Gbytes (2^{32}) of physical memory
 - Support for big-endian and true little-endian memory on a per-page basis
- Power management
 - Low power, 1.2-V design
 - Dynamic power management on the core minimizes power consumption of functional units, such as execution units, caches, and MMUs, when they are idle.
 - Core power-saving modes: core-halted and core-stopped
 - NAP, DOZE, and SLEEP bits in HID0 that can be used to assert *nap*, *doze*, and *sleep* core output signals to initiate power-saving modes at the integrated-device level
 - Internal clock multipliers of 2x, 2.5x, and 3x from bus clock
- Testability
 - LSSD scan design
 - JTAG interface
 - ESP support
 - ABIST for arrays
 - LBIST
- Reliability and serviceability
 - Internal code parity
 - Parity checking on e500 local bus

3.2 On-Chip Memory Unit

The MPC8540 contains an internal 256-Kbyte memory array that can be configured as memory-mapped SRAM or as a look-aside L2 cache. The array can also be divided into two 128-Kbyte arrays, one of which may be used as cache and the other as SRAM.

The memory controller for this array connects to the core complex bus (CCB) and communicates through 128-bit read and write buses to the e500 core and the MPC8540 system logic.

The on-chip memory unit contains:

- 256 Kbytes of on-chip memory
 - L2 cache partitioning is configurable
 - Can act as a 256-Kbyte L2 cache
 - 256-Kbyte array organized as 1024 eight-way sets of 32-byte cache lines
 - Array can be partitioned into 128-Kbyte L2 cache and 128-Kbyte memory mapped SRAM
 - Can act as two 128-Kbyte memory-mapped SRAM arrays or a 256-Kbyte SRAM region
 - SRAM operation is byte-accessible
 - Data ECC on 64-bit boundaries (single-error correction, double-error detection)
 - Tag parity (1 bit covering all tag bits)
 - Cache mode supports instruction caching, data caching, or both
 - External masters can force data to be allocated into the cache through programmed memory ranges or special transaction types
 - Separate locking for instructions and data so that locks can be set and cleared separately
 - Supports locking the entire cache or selected lines
 - Individual line locks are set and cleared through core-initiated instructions, by external reads or writes, or by accesses to programmed memory ranges
 - Flash clearing done through writes to L2 configuration registers
 - Locks for the entire cache may be set and cleared by accesses to memory-mapped control registers

3.2.1 On-Chip Memory as Memory-Mapped SRAM

When the on-chip memory is configured as an SRAM, the 256 Kbytes of memory can be configured to reside at any aligned location in the memory map. It is byte-accessible and fully ECC-protected, using read-modify-write transactions for sub-cacheline transactions. I/O devices can access the SRAM by marking transactions global so that they are directed to the CCB.

3.2.2 On-Chip Memory as L2 Cache

The MPC8540 on-chip memory arrays include a 256-Kbyte data array, an address tag array, and a status array.

The data array is organized as 1024 sets of 8 cache lines. Each cache line size is 32 bytes. The replacement policy in each eight-way set is governed by a pseudo-LRU algorithm. The data is protected with ECC, and the tag array is protected by parity.

The L2 cache tags are non-blocking for efficient load/store and snooping operations. The L2 cache can be accessed internally while a load miss is pending (allowing hits under misses). Subsequent to a load miss updating the memory, loads or stores can occur to that line on the very next cycle.

The L2 status array maintains status bits for each line to determine the status of the line. Different combinations of these bits result in different L2 states. Note that because the cache is always write-through, there is no modified state. The status bits include the following:

- V—Valid
- IL—Instruction locked
- DL—Data locked

All accesses to the L2 memory are fully pipelined so back-to-back loads and stores can have single-cycle throughput.

The cache can be configured to allocate instructions-only, data-only, or both. It can also be configured to allocate global I/O writes that correspond to a programmable address window or that use a special transaction type (stashing). In this way, DMA engines or I/O devices can force data into the cache.

Line locks can be set in a variety of ways. The Book E architecture defines instructions that explicitly set and clear locks in the L2. These instructions are supported by the core complex and the L2 controller. In addition, the L2 controller can be configured to lock all lines that fall into either of two specified address ranges when the line is allocated. Finally, the entire cache can be locked by writing to a configuration register in the L2 cache controller.

The status array tracks line locks as either instruction locks or data locks for each line, and the status array supports flash clearing of all instruction locks or data locks separately by writes to configuration registers in the L2 controller.

3.3 e500 Coherency Module (ECM)

The e500 coherency module (ECM) provides a mechanism for I/O-initiated transactions to snoop the bus between the e500 core and the integrated L2 cache in order to maintain coherency across local cacheable memory. It also provides a flexible switch-type structure for core and I/O-initiated transactions to be routed or dispatched to target modules on the device.

3.4 DDR SDRAM Controller

The MPC8540 supports DDR-I SDRAM that operates at up to 166 MHz (333-MHz data rate). The memory interface controls main memory accesses and provides for a maximum of 3.5 Gbytes of main memory. The memory controller can be configured to support the various memory sizes through software initialization of on-chip configuration registers.

The MPC8540 supports a variety of SDRAM configurations. SDRAM banks can be built using DIMMs or directly-attached memory devices. Fifteen multiplexed address signals provide for device densities of 64 Mbits, 128 Mbits, 256 Mbits, and 512 Mbits, and 1 Gbit. Four chip select signals support up to four banks of memory. The MPC8540 supports bank sizes from 64 Mbytes to 1 Gbyte. Nine-column address strobes (MDM[0:8]) are used to provide byte selection for memory bank writes.

The MPC8540 can be configured to retain the currently active SDRAM page for pipelined burst accesses. Page mode support of up to 16 simultaneously open pages can dramatically reduce access latencies for page hits. Depending on the memory system design and timing parameters, using page mode can save 3 to 4 clock cycles from subsequent burst accesses that hit in an active page.

The MPC8540 supports ECC for system memory. Using ECC, the MPC8540 detects and corrects all single-bit errors and detects all double-bit errors and all errors within a nibble.

The MPC8540 can invoke a level of system power management by asserting the MCKE SDRAM signal on-the-fly to put the memory into a low-power sleep mode.

3.5 Programmable Interrupt Controller (PIC)

The programmable interrupt controller (PIC) implements the necessary functions to provide a flexible solution for a general-purpose interrupt control. The interrupt controller unit implements the logic and programming structures of the OpenPIC architecture. The MPC8540 interrupt controller unit supports its processor core and provides for 12 external interrupts (with fully nested interrupt delivery), 4 message interrupts, internal-logic driven interrupts, and 4 global high resolution timers. Up to 16 programmable interrupt priority levels are supported.

The interrupt controller unit can be bypassed to allow use of an external interrupt controller. Inter-processor interrupt (IPI) communication is supported through the external interrupt and core reset signals of different processor cores on the same device. The four IPIs are only used for self-interrupt in a single-core device such as the MPC8540.

3.6 I²C Controller

The inter-IC (IIC or I²C) bus is a two-wire, bidirectional serial bus that provides a simple and efficient method of data exchange between devices. The synchronous, multiple master bus of the I²C allows the MPC8540 to exchange data with other I²C devices, such as microcontrollers, EEPROMs, real-time clock devices, A/D converters, and LCDs. The two-wire bus (serial data SDA and serial clock SCL) minimizes the interconnections between devices. The synchronous, multiple master bus of the I²C allows the connection of additional devices to the bus for expansion and system development.

The I²C controller is a true multiple master bus; it includes collision detection and arbitration that prevents data corruption if two or more masters attempt to control the bus simultaneously. This feature allows for complex applications with multiprocessor control. The I²C controller consists of a transmitter/receiver unit, a clocking unit, and a control unit. The I²C unit supports general broadcast mode, and on-chip filtering rejects spikes on the bus.

3.7 Boot Sequencer

The MPC8540 provides a boot sequencer that uses the I²C interface to access an external serial ROM and loads the data into the MPC8540's configuration registers. The boot sequencer is enabled by a configuration pin sampled at the negation of the MPC8540 hardware reset signal. If enabled, the boot sequencer holds the MPC8540 processor core in reset until the boot sequence is complete. If the boot sequencer is not enabled, the processor core exits reset and fetches boot code in default configurations.

3.8 Dual Universal Asynchronous Receiver/Transmitter (DUART)

The MPC8540 includes a DUART intended for use in maintenance, bringing-up, and debugging of systems. The MPC8540 provides a standard four-wire handshake (SIN, SOUT, RTS, CTS) for each port. The DUART is a slave interface. An interrupt is provided to the interrupt controller or optionally steered externally to allow device handshakes. Interrupts are generated for transmit, receive, line status, and MODEM status.

The MPC8540 DUART supports full-duplex operation. It is compatible with the PC16450 and PC16550 programming models. Also, 16-byte FIFOs are supported for both the transmitter and the receiver.

Software programmable baud generators divide the system clock to generate a 16x clock. Serial interface data formats (data length, parity, 1/1.5/2 STOP bit, baud rate) are also software selectable.

3.9 10/100 Fast Ethernet Controller

The MPC8540 provides a 10/100 fast Ethernet controller operating at 10 to 100 Mbps as a device debug and maintenance port. The gigabit three-speed Ethernet controllers (TSECs) are separate from this 10/100 fast Ethernet controller and are defined and described in Chapter 13, “Three-Speed Ethernet Controllers,” in the *MPC8540 Integrated Host Processor Reference Manual*.

The 10/100 Ethernet standard supports the IEEE 802.3 Ethernet frame format, backward compatibility for installed media, and the use of full- or half-duplex CSMA/CD. The Ethernet protocol implements the bottom two layers of the open systems interconnection (OSI) 7-layer model, that is, the data-link and physical sublayers.

3.10 Local Bus Controller (LBC)

The MPC8540 local bus controller (LBC) port allows connections with a wide variety of external memories, DSPs, and ASICs. Three separate state machines share the same external pins and can be programmed separately to access different types of devices. The general-purpose chip select machine (GPCM) controls accesses to asynchronous devices using a simple handshake protocol. The user programmable machine (UPM) can be programmed to interface to synchronous devices or custom ASIC interfaces. The SDRAM controller provides access to standard SDRAM. Each chip select can be configured so that the associated chip interface can be controlled by the GPCM, UPM, or SDRAM controller. All may exist in the same system.

The GPCM provides a flexible asynchronous interface to SRAM, EPROM, FEPRM, ROM, and other devices such as asynchronous DSP host interfaces and CAMs. Minimal glue logic is required. Handshake signals can be configured to transition on fractions of the system clock. The GPCM does not support bursting.

The UPM allows an extremely flexible interface in which the programmer configures each of a set of general-purpose protocol signals by writing the transition pattern into a memory array. The UPM supports synchronous and bursting interfaces. It also supports multiplexed addressing so that a simple DRAM interface can be implemented. The UPM is entirely flexible in order to provide a very high degree of customization with respect to both asynchronous and burst-synchronous interfaces, which permits glueless or almost glueless connection to burst SRAM, custom ASIC, and synchronous DSP interfaces.

The LBC provides a synchronous DRAM (SDRAM) machine that supplies the control functions and signals for glueless connection to JEDEC-compliant SDRAM devices. An internal DLL (delay-locked loop) for bus clock generation ensures improved data setup margins for board designs. The SDRAM machine can optimize burst transfers and exploits interleaving to maximize data transfer bandwidth and minimize access latency. Programmable row and column address multiplexing allows a variety of SDRAM configurations and sizes to be supported without hardware changes.

3.11 Three-Speed Ethernet Controllers (10/100/1Gb)

The MPC8540 has two on-chip three-speed Ethernet controllers (TSECs). The TSECs incorporate a media access control sublayer (MAC) that supports 10- and 100-Mbps, and 1-Gbps Ethernet/802.3 networks with MII, GMII, RGMII, RTBI, and TBI physical interfaces. The TSECs include 2-Kbyte receive and transmit FIFOs, and DMA functions.

The buffer descriptors are based on the MPC8260 and MPC860T 10/100 programming models.

The MPC8540 TSECs support programmable CRC generation and checking, RMON statistics, and jumbo frames of up to 9.6 Kbytes. Frame headers and buffer descriptors can be forced into the L2 cache to speed classification or other frame processing.

3.12 Integrated DMA

The MPC8540 DMA engine is capable of transferring blocks of data from any legal address range to any other legal address range. Therefore, it can perform a DMA transfer between any of its I/O or memory ports or even between two devices or locations on the same port.

The four-channel DMA controller allows chaining (both extended and direct) through local memory-mapped chain descriptors. Scattering, gathering, and misaligned transfers are supported. In addition, advanced capabilities such as stride transfers and complex transaction chaining are supported.

DMA transfers can be initiated by a single write to a configuration register. There is also support for external control of transfers using `DMA_DREQ`, `DMA_DACK`, and `DMA_DDONE` handshake signals.

DMA descriptors encompass a rich set of attributes that allow DMA transfers to bypass outbound address translation and supply external addresses and attributes directly to the RapidIO port. Local attributes such as snoop and L2-write stashing can be specified by descriptors.

Interrupts are provided on a completed segment, link, list, chain, or on an error condition. Coherency is selectable and hardware enforced (snoop/no snoop).

3.13 PCI Controller

The MPC8540 64-bit PCI controller is compatible with the *PCI Local Bus Specification, Revision 2.2* and the *PCI-X Addendum, Revision 1.0*. The interface can function as a host or agent bridge interface in either PCI or PCI-X mode. Both PCI and PCI-X modes support 64-bit addressing and 32-bit or 64-bit data buses.

As a master, the MPC8540 supports read and write operations to the PCI memory space, the PCI I/O space, and the PCI configuration space. Also, the MPC8540 can generate PCI special-cycle and interrupt-acknowledge commands. As a target, the MPC8540 supports read and write operations to system memory as well as configuration accesses.

PCI-X functionality includes split transaction support for four outstanding split transactions. Split response data is returned in order without interleaving. As a target, the MPC8540 supports all PCI-X sizes. As a master it internally combines transactions up to 256 bytes.

An internal arbiter can be used to support up to five external masters. A round robin arbitration algorithm with two priority levels is used.

3.14 RapidIO Controller

The RapidIO interconnect unit on the MPC8540 is based on the *RapidIO Interconnect Specification, Revision 1.1*. RapidIO is a high-performance, point-to-point, low-pin-count, packet-switched system-level interconnect that can be used in a variety of applications as an open standard. The RapidIO architecture provides a rich variety of features including high data bandwidth, low-latency capability, and support for high-performance I/O devices, as well as providing message-passing and software-managed programming models.

The RapidIO unit on the MPC8540 supports the I/O and message-passing logical specifications, the common transport specification, and the 8/16 LP-LVDS physical layer specification of the *RapidIO Interconnect Specification*. It does not support the globally shared memory logical specification.

Highlights of the implementation include: support for four priority levels and ordering within a priority level, CRC error management, 32- to 256-byte transactions and 8-bit data width ports.

The physical layer of the RapidIO unit can operate at up to 500 MHz. Because the interface is defined as a source-synchronous, double-data-rate, LVDS-signaling interconnect, the theoretical unidirectional peak bandwidth

is 8 Gbps. Receive and transmit ports operate independently, resulting in an aggregate theoretical bandwidth of 16 Gbps.

3.15 RapidIO Message Unit

The MPC8540's RapidIO messaging supports one inbox/outbox structure for data and one doorbell structure for messages. Both chaining and direct modes are provided for the outbox, and messages can hold up to 16 packets of 256 bytes, or a total of 4 Kbytes.

3.16 Power Management

In addition to low-voltage operation and dynamic power management in its execution units, the MPC8540 supports four power consumption modes: full-on, doze, nap, and sleep. The three low-power modes: doze, nap, and sleep, can be entered under software control in the e500 core or by external masters accessing a configuration register.

Doze mode suspends execution of instructions in the e500 core. The core is left in a standby mode in which cache snooping and time base interrupts are still enabled. Device logic external to the processor core is fully functional in this mode.

Nap mode shuts down clocks to all the e500 functional units except the time base, which can be disabled separately. No snooping is performed in nap mode, but the device logic external to the processor core is fully functional.

Sleep mode shuts down not only the e500 core, but all of the MPC8540 I/O interfaces as well. Only the interrupt controller and power management logic remain enabled so that the device can be awakened.

3.17 Clocking

The MPC8540 takes in the PCI_CLK/SYSCLK signal as an input to the device PLL and multiplies it by an integer from 1 to 16 to generate the core complex bus clock (the platform clock), which operates at the same frequency as the DDR DRAM data rate (for example, 266 or 333 MHz). The L2 cache also operates at this frequency. The e500 core uses the CCB clock as an input to its PLL, which multiplies it again by 2, 2.5, 3, or 3.5 to generate the core clock.

DLLs are used in the DDR SDRAM controller and the local bus memory controller (LBC) to generate memory clocks. Six differential clock pairs are generated for DDR SDRAMs. Two clock outputs are generated for the LBC.

The RapidIO transmit clock may be sourced from one of three locations: the platform clock, the RapidIO receive clock, or a special differential clock input. This input is designed to receive inputs from an external clock synthesis device driving a clock with a frequency of up to 500 MHz.

3.18 Address Map

The MPC8540 supports a flexible physical address map. Conceptually, the address map consists of local space and external address space. The local address map is 4 Gbytes. The MPC8540 can be made part of a larger system address space through the mapping of translation windows. This functionality is included in the address translation and mapping units (ATMUs). Both inbound and outbound translation windows are provided. The ATMUs allows the MPC8540 to be part of larger address maps such as the PCI 64-bit address environment and the RapidIO environment.

3.19 OCeaN Switch Fabric

In order to reduce the strain on the core interconnects with the addition of new functional blocks in this generation of the PowerQUICC family, an on-chip non-blocking crossbar switch fabric called OCeaN (on-chip network) has been integrated to decrease contention, decrease latency, and increase bandwidth. This revolutionary non-blocking crossbar fabric allows for full-duplex port connections at 128 Gbps concurrent throughput and independent per-port transaction queuing and flow control.

3.20 Processing Across the On-Chip Fabric

When processing across the on-chip fabric, the ATMUs at each fabric port are used to determine the flow of data across the MPC8540. The ATMUs at each fabric port are responsible for generating a fabric port destination ID as well as a new local device address. The port ID and local address are based on the programmed destination of the transaction. The following is a general overview of how the ATMUs process transactions over the on-chip fabric. (Refer to [Figure 2](#).)

1. When a transaction on one of the fabric ports begins, the ATMU on the origination port translates the programmed destination address into both a destination fabric port ID and a local device address.
2. The data is then processed across the on-chip fabric from the origination port to the destination port.
3. If the destination port connects off-chip (for example, to a PCI or RapidIO device), the local device address is translated by the destination port ATMU to an outbound address with respect to the destination port's memory map, and the data is processed accordingly.




Figure 2. Processing Transactions Across the On-Chip Fabric

3.21 Data Processing with the e500 Coherency Module

Processing through the ECM is similar to processing across the on-chip fabric (in the sense of how data is received and transmitted) with the exception that the transaction passes through the ECM. The purpose of the ECM is to provide a means for any I/O transaction to maintain coherency with the cacheable DDR SDRAM and the local bus memory. However, simply using the ECM does not make transactions across it coherent. The e500 and L2 cache are snooped to maintain coherency only if the transaction across the ECM is designated as global (GBL bit set). Otherwise, the transaction passes through the ECM using the ECM as a simple conduit to get to its destination. In essence, only global transactions across the ECM are coherent transactions; all other transaction (across the on-chip fabric) are non-coherent.

4 MPC8540 Application Examples

The following section provides block diagrams of different MPC8540 applications. The MPC8540 is a very flexible device and can be configured to meet many system application needs. In order to build a system, many factors should be considered.

4.1 MPC8540 Applications

The MPC8540 can be used for control processing in applications such as routers, switches, network memory, internet access devices, storage systems, and image display systems (IADs). Figure 3 shows an MPC8540 as a control plane processor in a high-performance communication system using distributed processing.

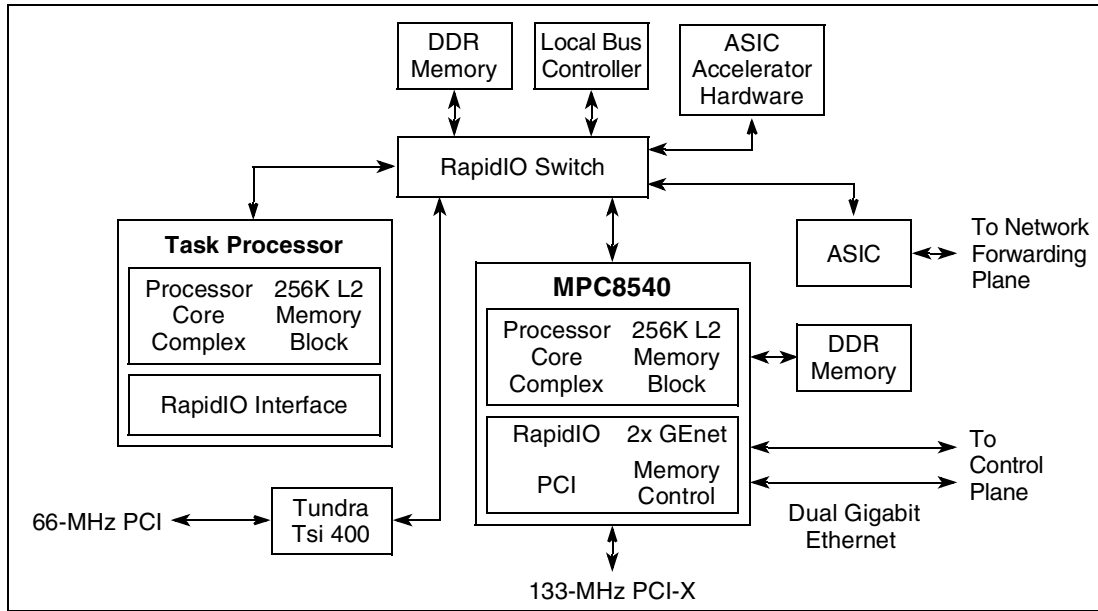


Figure 3. High-Performance Communication System Using Distributed Processing

Figure 4 shows the MPC8540 as part of a high-end network card used in a system area network that is enabled by PCI/PCI-X.

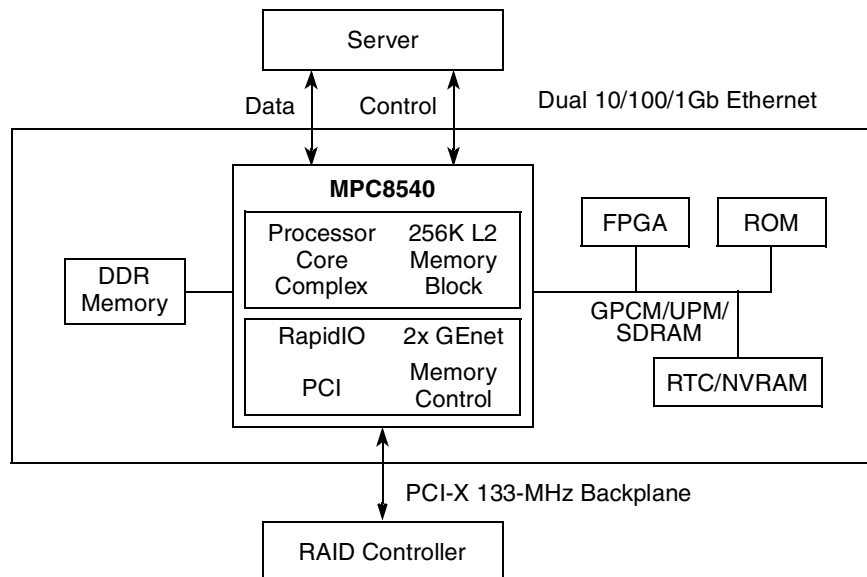


Figure 4. High-Performance Communication System Using MPC8540

Figure 5 shows the MPC8540 in a redundant array of independent disks (RAID) controller application.

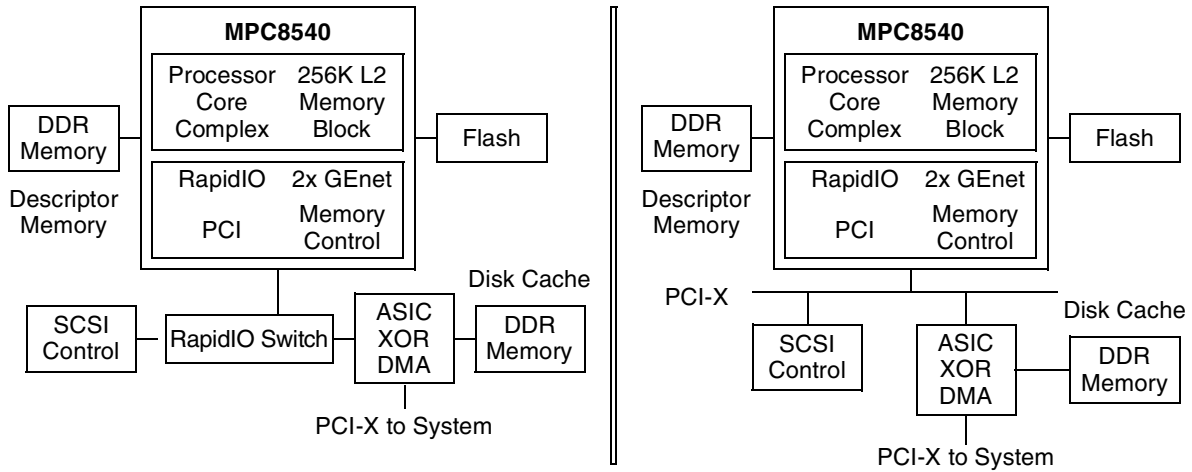


Figure 5. RAID Controller Application Using MPC8540

Figure 6 illustrates the MPC8540 in a virtual private network (VPN) access that is enabled through RapidIO and Ethernet.




Figure 6. VPN Access Point Enabled by RapidIO and Ethernet

Revision History

Figure 7 shows the MPC8540 in a serializer/deserializer (SERDES) data transfer mode application enabled by Ethernet.

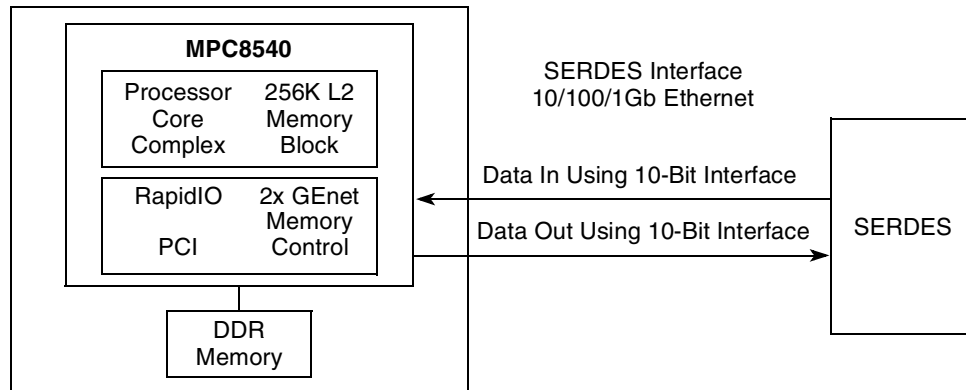


Figure 7. MPC8540 with SERDES

Figure 8 shows a DSP farm enabled by the MPC8540 using its local bus controller port.

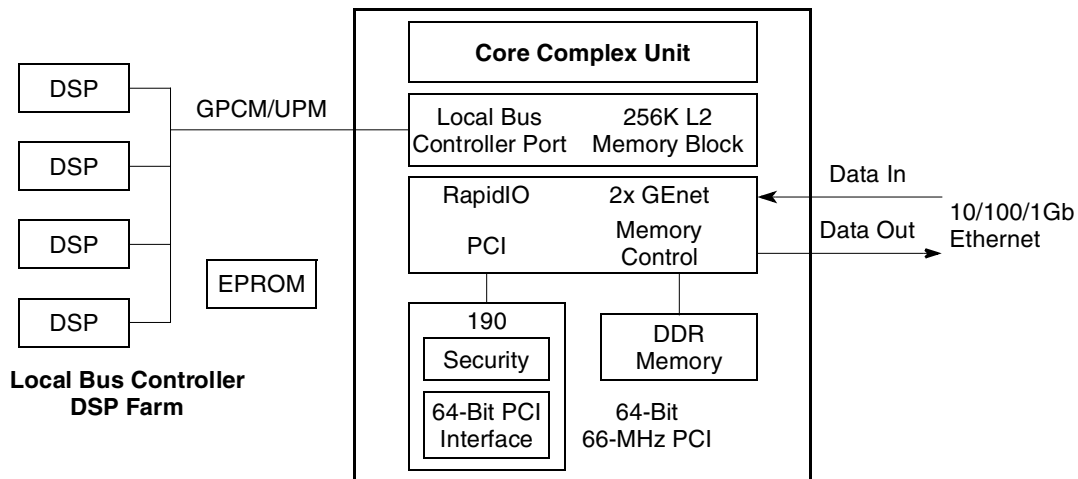


Figure 8. MPC8540 with a DSP Farm Enabled by General-Purpose Chip Select Machine or User Programmable Machine on the Local Bus Controller Port

5 Revision History

Table 1 provides a revision history for this template.

Table 1. Document Revision History

Revision	Date	Substantive Change(s)
0.1	1/12/2005	Replaced "symmetric multiprocessing(SMP)" with "distributed processing."
0	12/03/2003	Initial release.

THIS PAGE INTENTIONALLY LEFT BLANK

THIS PAGE INTENTIONALLY LEFT BLANK

THIS PAGE INTENTIONALLY LEFT BLANK

How to Reach Us:

Home Page:

www.freescale.com

email:

support@freescale.com

USA/Europe or Locations Not Listed:

Freescale Semiconductor
Technical Information Center, CH370
1300 N. Alma School Road
Chandler, Arizona 85224
(800) 521-6274
480-768-2130
support@freescale.com

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)
support@freescale.com

Japan:

Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku
Tokyo 153-0064, Japan
0120 191014
+81 2666 8080
support.japan@freescale.com

Asia/Pacific:

Freescale Semiconductor Hong Kong Ltd.
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate,
Tai Po, N.T., Hong Kong
+800 2666 8080
support.asia@freescale.com

For Literature Requests Only:

Freescale Semiconductor
Literature Distribution Center
P.O. Box 5405
Denver, Colorado 80217
(800) 441-2447
303-675-2140
Fax: 303-675-2150
LDCForFreescaleSemiconductor
@hibbertgroup.com

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part.

Freescale™ and the Freescale logo are trademarks of Freescale Semiconductor, Inc. The PowerPC name is a trademark of IBM Corp. and is used under license. All other product or service names are the property of their respective owners.

© Freescale Semiconductor, Inc. 2004.

MPC8540PB
Rev. 0.1
1/2005

