Technical Summary

MPC8560TS/D
Rev. 0.6 7/2002

MPC8560
Power QUICC III™ Integrated Communications Processor

MOTOROLA intelligence everywhere™ | digital dna™

The MPC8560 PowerQUICC III is a next-generation PowerQUICC II integrated communications processor. The MPC8560 provides integration of processing power for networking and communications peripherals resulting in higher device performance.The MPC8560 is a member of the growing family of products that combine system level support for industry standard interfaces to processors that implement the PowerPC architecture.

One block of the MPC8560 is a high-performance embedded e500 core processor. The e500 core, with its 256 Kbyte of level-2 cache, implements the enhanced Book E instruction set architecture and provides unprecedented levels of hardware and software debug support.

A second block is the communications processor module (CPM). The CPM of the MPC8560 supports three high-performance fast communications channels (FCCs) for 155 Mbps ATM and Fast Ethernet, and up to 256 full-duplex, time-division-multiplexed (TDM) channels using two multi-channel controllers (MCCs).

In addition, the MPC8560 offers two integrated 10/100 Mbps/1 Gbps triple speed Ethernet controllers (TSECs), a DDR SDRAM memory controller, a 64-bit PCI/PCI-X controller, and an 8-bit RapidIO port, a programmable interrupt controller, an I2C controller, a 4-channel DMA controller, and a general purpose I/O port. The MPC8560 is a member of the growing family of products that combine system level support for industry standard interfaces to processors that implement the PowerPC architecture. This high level of integration simplifies board design and offers significant bandwidth and performance for high-end control-plane and data-plane applications.

This document describes the functional operation of the MPC8560 with emphasis on peripheral functions. This document is divided into four parts:

- Part I, “MPC8560 Overview,” provides an high-level overview of the MPC8560 features, including a block diagram showing the major components.
- Part II, “MPC8560 Architecture Overview,” describes the major functional units of the MPC8560 and its interfaces.
- Part III, “MPC8560 Application Examples,” provides block diagrams of different MPC8560 applications. The MPC8560 is a very flexible device and can be configured to meet many system application needs.
- Part IV, “Compatibility Issues,” describes some software and hardware compatibility issues as well as differences between the MPC8560 and the MPC8260.

Part I MPC8560 Overview

The following section provides an high-level overview of the MPC8560 features.

Figure 1-1 shows the major functional units within the MPC8560.

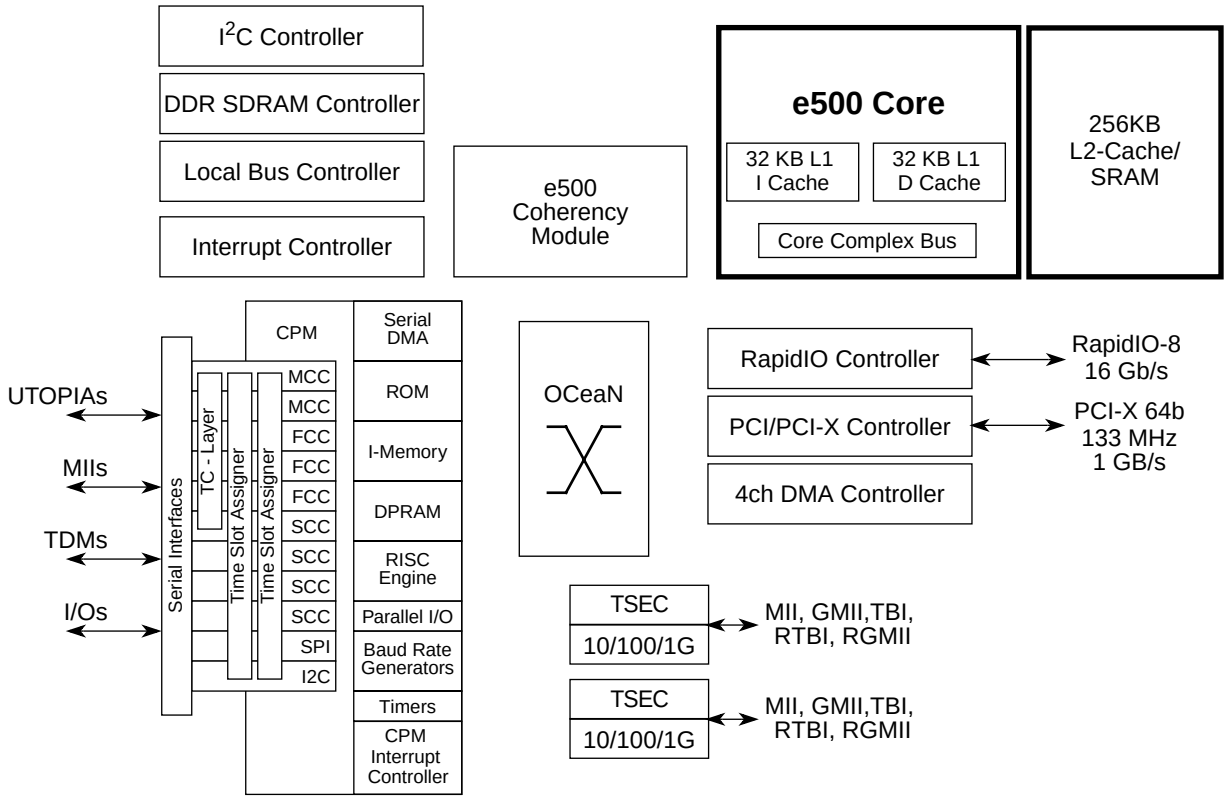

Figure 1-1. MPC8560 **Block Diagram**

The following lists an overview of the MPC8560 feature set.

- Embedded e500 Book E-compatible core
  - — High-performance, 32-bit Book E enhanced core that implements the PowerPC architecture
  - — Dual issue superscalar, 7-stage pipeline design
  - — 32-Kbyte L1 instruction cache and 32-Kbyte L1 data cache with parity protection
  - — Lockable L1 caches—entire cache or on a per-line basis
  - — Separate locking for instructions and data
  - — Single processing engine auxiliary processing unit (SPE APU)

— Single-precision floating-point single-instruction, multiple-data (SIMD) operations.
— Memory management unit especially designed for embedded applications
— Enhanced hardware and software debug support
— Dynamic power management
— Performance monitor facility

- High-performance RISC CPM operating at up to 333 MHz
  — CPM software compatibility with previous PowerQUICC families
  — One instruction per clock
  — Executes code from internal ROM or instruction RAM
  — 32-bit RISC architecture
  — Tuned for communication environments: instruction set supports CRC computation and bit manipulation.
  — Internal timer
  — Interfaces with the embedded e500 core processor through a 32-Kbyte dual-port RAM and virtual DMA channels for each peripheral controller
  — Handles serial protocols and virtual DMA.
  — Three full-duplex fast serial communications controllers (FCCs) support the following protocols:
    – ATM protocol through UTOPIA interface (FCC1 and FCC2 only)
    – IEEE802.3/Fast Ethernet
    – HDLC
    – Totally transparent operation
  — Two multi-channel controllers (MCCs) that together can handle up to 256 HDLC/transparent channels at 64 Kbps each, multiplexed on up to eight TDM interfaces
  — Four full-duplex serial communications controllers (SCCs) support the following protocols:
    – High level/synchronous data link control (HDLC/SDLC)
    – LocalTalk (HDLC-based local area network protocol)
    – Universal asynchronous receiver transmitter (UART)
    – Synchronous UART (1x clock mode)
    – Binary synchronous communication (BISYNC)
    – Totally transparent operation
  — Serial peripheral interface (SPI) support for master or slave

— I^2C bus controller
— Time-slot assigner supports multiplexing of data from any of the SCCs and FCCs onto eight time-division multiplexed (TDM) interfaces. The time-slot assigner supports the following TDM formats:
  – T1/CEPT lines
  – T3/E3
  – Pulse code modulation (PCM) highway interface
  – ISDN primary rate
  – Motorola interchip digital link (IDL)
  – General circuit interface (GCI)
— User-defined interfaces
— Eight independent baud rate generators (BRGs)
— Four general-purpose 16-bit timers or two 32-bit timers
— General-purpose parallel ports—16 parallel I/O lines with interrupt capability

- Address translation and mapping unit (ATMU)
  — 4 inbound windows per port
  — 8 outbound windows per port
  — 12 local target mapping windows
- DDR memory controller
  — Support for DDR SDRAM
  — 166 MHz clock, 333 MHz data rate
  — 64-bit dedicated data bus
  — Programmable settings for meeting all SDRAM timing parameters
  — Many different SDRAM configurations supported
  — Full ECC support and error injection
  — Support for data mask signals and read-modify-write for sub-double word writes
  — Support for double-bit error detection and single-bit error correction ECC (8-bit check word across 64-bit data)
  — Two-entry input request queue
  — Open page management (dedicated entry for each sub-bank)
  — Memory controller clock frequency of two times the SDRAM clock with support for sleep power management mode
    – Support for error injection
  — Sleep mode support for self refresh SDRAM, supports auto refreshing, and on-the-fly power management using CKE

  - Page mode support (up to eight simultaneous open pages)
  - Contiguous or discontiguous memory mapping
  - Read-modify-write support for Atomic Inc, Dec, Set, Clear
  - Registered DIMM support
  - Fast memory access via JTAG port
  - 2.5V SSTL2 compatible I/O
- Interrupt controller
  - Supports 16 programmable interrupt and processor task priority levels
  - Supports 12 discrete external interrupts
  - Supports four message interrupts
  - Supports connection of external interrupt controller
  - Fully nested interrupt delivery
  - Four global high resolution timers/counters
  - Processor initialization control and soft reset
  - Software reset of the interrupt controller
- I^2C unit
  - Two-wire interface
  - Multiple master support
  - Arbitration lost interrupt with automatic mode switching from master to slave
  - Calling address identification interrupt
  - START and STOP signal generation/detection
  - Acknowledge bit generation/detection
  - Bus busy detection
  - Software-programmable clock frequency
  - Software-selectable acknowledge bit
  - On-chip filtering for spikes on the bus
- Boot sequencer
  - Loads configuration registers from serial ROM at reset via I^2C interface
- Local Bus Controller
  - Multiplexed 32-bit address and data operating up to 167 MHz
  - Single-master bus, supports external slaves
  - Up to eight-beat burst transfers
  - 32-, 16-, and 8-bit port sizes controlled by on-chip memory controller
  - Three user programmable machines (UPM)

— Dedicated single data rate SDRAM controller
— General purpose chip select machine (GPCM)
— Eight chip selects
— Parity support
— Boot chip select with programmable bus width (8-,16- or 32-bit)
— SDRAM machine

- Triple speed (10/100/1000) Ethernet controllers (TSEC)
  — Dual IEEE 802.3, 802.3u, 820.3x, 802.3z, 802.3 AC compliant controllers
  — Support for different Ethernet physical interfaces:
    – 10/100/1000 Mbps IEEE 802.3 GMII
    – 10/100 Mbps IEEE 802.3 MII
    – 10 Mbps IEEE 802.3 MII
    – 1000 Mbps IEEE 802.3z TBI
    – 10/100/1000 Mbps RGMII/RTBI
  — Full and half-duplex support
  — Buffer descriptors are backwards compatible with MPC8260 and MPC860T 10/100 programming models
  — Layer 2 Acceleration
    – 8 unicast exact address matches
    – Up to 16 4-byte pattern matches
    – Accept or reject on address or pattern match
    – 256 entry hash for unicast and multicast
    – Direct queuing of four flows
    – Packet field extraction and insertion
  — 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

- 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 up to 256 bytes
- Integrated DMA controller
  - 4-channel controller
  - All channels accessible by both the local and remote masters
  - Extended DMA operations (advanced chaining and stride capability)
  - Support for cascading descriptor chains
  - Extended chaining and direct modes
  - Scatter gathering
  - Unaligned transfer capability
  - Interrupt on completed segment, link, list, and error
  - Supports transfers to or from any local memory or I/O port
  - Supports transfers of packet data using packet buffer descriptors
  - Selectable hardware-enforced coherency (snoop/no-snoop)
  - Ability to start DMA from external 3-pin interface
  - Programmable chunk wire
  - Ability to launch DMA from single write transaction
  - Optimal 3-pin hardware handshake interface per channel
- PCI/PCI-X functional unit
  - PCI 2.2 and PCI-X 1.0 compatible
  - 64-bit or 32-bit PCI support at 16 MHz 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
  - Support posting of processor to PCI and PCI to memory writes
  - PCI 3.3V compatible
  - Selectable hardware-enforced coherency

- 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, Clear
  - LVDS signalling
- 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
- Power management
  - Fully static 1.2V CMOS design with 3.3V and 2.5 V I/O
  - Supports power save modes: nap, doze, and sleep
  - Supports suspend power save mode
  - Supports dynamic power management
- System performance monitor
  - Support eight 32-bit counters that count the occurrence of selected events
  - Supports 64 reference events which can be counted on any of the eight counters
  - Supports duration and quantity threshold counting
  - Burstiness feature that permits counting of burst events with a programmable time between bursts

- 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
- Supports inverse muxing of ATM cells (IMA)

Part II MPC8560 Architecture Overview

The following section describe the major functional units of the MPC8560.

2.1 e500 Core Overview

The MPC8560 uses the e500 microprocessor core complex. Both the e500 core and the CPM have an internal PLL, that allows independent optimization of their operating frequencies. The core and CPM frequencies are derived from either the primary PCI clock input or an external oscillator. For information regarding the e500 core refer to the following specifications:

- *EREF: A Reference for Motorola Book E and the e500 Core*

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

- Implements full Book E 32-bit architecture
- Cache structure
  - 32–Kbyte, 32-byte line, 8-way set associative instruction cache
  - 32–Kbyte, 32-byte line, 8-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
  - Supports all Book E memory coherency modes*
- 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 (one 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
- Dual-issue superscalar control
  - Two-instructions-per-clock peak issue rate to any two distinct EUs
  - Precise exception handling
- 7-stage pipeline control
- Decode unit
  - 12-entry instruction queue
  - Full hardware detection of interlocks
  - Decode up to two instructions per cycle
  - Decode serialization control
  - Register dependency resolution and renaming
- Branch processing unit (BPU)
  - Dynamic branch prediction
  - Two-entry branch instruction queue
  - 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 units
  - Single-cycle add and subtract
  - Single-cycle shift and rotate
  - Single-cycle logical operations
  - Supports integer signal processing operations

- Multiple-cycle unit (MU)
  - Four-cycle latency for integer and SPE multiplication (including SPE vector and scalar, integer, fractional, and floating-point multiply instructions).
  - Variable-latency divide: 4, 11, 19, and 35 cycles for all Book E and SPE divide instructions. Note that the MU allows divide instructions to bypass the first two MU pipeline stages, freeing those stages for other MU instructions to execute in parallel. Two cycles before the divide instruction needs stage three, the MU signals the GIQ not to issue instructions to it, making stage three (MU-3) available for the divide instruction to reenter the MU pipeline.
  - Four-cycle floating-point multiply
  - Four-cycle floating-point add and subtract
- Signal processing engine APU (SPE APU). The SIMD capability provided by the double wide 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
  - 4-cycle integer pipelined multiplies
  - 4-, 11-, 19-, and 35-cycle integer divides
  - 4-cycle single instruction multiple data (SIMD) pipelined multiply-accumulate (MAC)
  - 64-bit accumulator for no-stall MAC operations
  - Single precision floating-point single instruction multiple data (SIMD) operations
- Load/store unit (LSU)
  - 3-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
  - Supports both three-state MEI cache coherency and four-state cache coherency: invalid, shared, exclusive, exclusive-modified (MESI)
  - 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, 1 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
    - – Vector floating-point data exception
    - – Vector floating-point round exception
    - – Performance monitor
- Memory management unit
  - — Data L1 MMU
    - – 4-entry, fully-associative TLB array for variable-sized pages
    - – 64-entry, 4-way set-associative TLB for 4-Kbyte pages
  - — Instruction L1 MMU
    - – 4-entry, fully-associative TLB array for variable-sized pages
    - – 64-entry, 4-way set-associative TLB for 4-Kbyte pages
  - — Unified L2 MMU
    - – 16-entry, fully-associative TLB array for variable-sized pages
    - – 256-entry, 2-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- and true little-endian memory on a per-page basis

- Power management
  — Low power design
  — 1.2 volt design
  — Internal clock multipliers of 2x, 2.5x, 3x from bus clock
  — Power-saving modes: standby and shutdown
  — Dynamic power management of execution units, caches, and MMUs
- Testability
  — LSSD scan design
  — JTAG interface
  — ESP support
  — ABIST for arrays
  — LBIST
- Reliability and serviceability
  — Internal code parity
  — Parity checking on e500 local bus

2.2 On-Chip Memory Unit

The MPC8560 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 be divided into two 128-Kbyte arrays, one of which may be used in cache mode.

The memory controller for this array connects to the core complex bus (CCB) and communicates via 128-bit read and write buses to the e500 core and the MPC8560 system logic.

2.2.1 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.

2.2.2 On-Chip Memory as L2 Cache

The MPC8560 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. Replacement policy within each 8-way set is governed by a pseudo-LRU algorithm.

The data is protected with ECC, while 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 and stores can occur to that line on the very next cycle.

The L2 status array stores five bits for each line that are used to determine the status of the line. Different combinations of these bits result in different L2 states. Note that for write-through accesses, there is no modified state. The status bits are:

V -> Valid

S -> Shared

IL -> Instruction Locked

DL -> Data Locked

T -> Stale

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. 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 cache controller.

2.3 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 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 chip.

2.4 Communications Processor Module (CPM)

The CPM contains features that allow the MPC8560 to excel in a variety of applications targeted for networking and telecommunication markets. The MPC8560 CPM is a superset of the MPC8260 PowerQUICC II, with enhanced CP performance. The CPM also has additional hardware and microcode routines that support high bit rate protocols like ATM (up to 155 Mbps full-duplex) and Fast Ethernet (100 Mbps full-duplex).

Figure 2-2 shows the major functional units within the MPC8560 CPM.

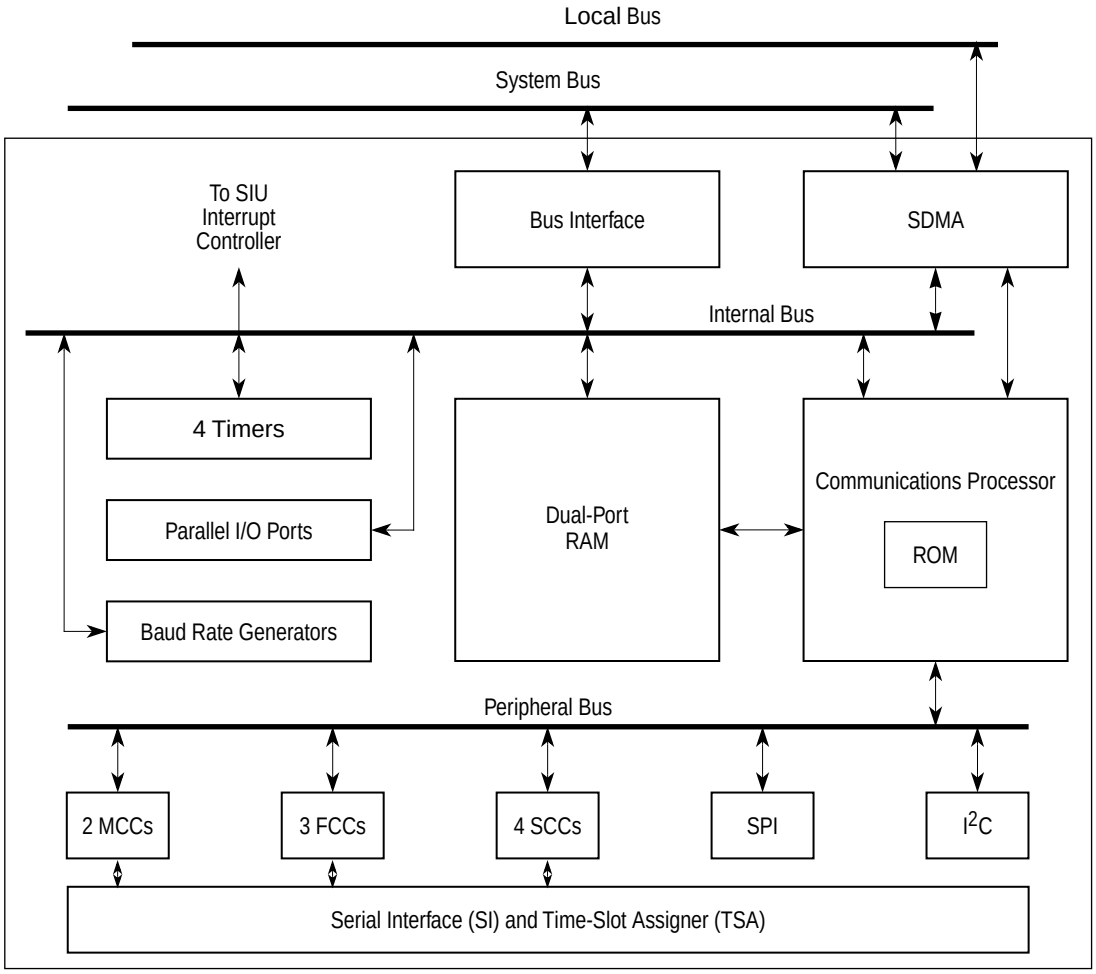

Figure 2-2. MPC8560 CPM Block Diagram

The following list summarizes the major features of the CPM:

- The communications processor (CP) is an embedded 32-bit RISC controller residing on a separate bus (CPM local bus). With this separate bus, the CP does not affect the performance of the e500 core. The CP handles the lower-layer tasks and DMA control activities, leaving the e500 core free to handle higher-layer activities. The CP has an instruction set optimized for communications but that can also be used for general-purpose applications, relieving the system core of small, often repeated tasks.
- Two serial DMAs (SDMAs), one associated with the local bus and one associated with the coherency module, handling transfers simultaneously
- Three full-duplex, serial fast communications controllers (FCCs) supporting ATM (155 Mbps) protocol through UTOPIA level 2 interface (there are two UTOPIA interfaces on the MPC8560), IEEE 802.3 and Fast Ethernet protocols, HDLC up to E3 rates (45 Mbps) and totally transparent operation. Each FCC can be configured to transmit fully transparent and receive HDLC data or vice-versa.
- Two multichannel controllers (MCCs) that can handle an aggregate of 256 x 64 Kbps HDLC or transparent channels, multiplexed on up to eight TDM interfaces. The MCC also supports super-channels of rates higher than 64 Kbps and subchanneling of the 64 Kbps channels.
- Four full-duplex serial communications controllers (SCCs) supporting, high-level synchronous data link control, HDLC, local talk, UART, synchronous UART, BISYNC, and transparent.
- Serial peripheral interface (SPI) and I^2C bus controllers
- Time-slot assigner (TSA) that supports multiplexing of data from any of the four SCCs and three FCCs
- ATM TC-layer functionality is implemented internally to support applications which receive ATM traffic over standard serial protocols (T1, E1, xDSL) via its serial interface ports

2.5 DDR SDRAM Controller

The MPC8560 supports DDR-I SDRAM 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 MPC8560 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, 512 Mbits, and 1 Gbits. Four chip select signals support up to four banks of memory. The MPC8560 supports bank

sizes from 64 Mbytes to 1 Gbyte. Nine column address strobes ($\overline{DM}$[0:8]) signals are used to provide byte selection for memory bank writes.

The MPC8560 can be configured to retain the currently active SDRAM page for pipelined burst accesses. Page mode support of up to eight 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 three to four clock cycles from subsequent burst accesses that hit in an active page.

The MPC8560 supports error checking and correction (ECC) for system memory. Using ECC, the MPC8560 detects and corrects all single-bit errors and detects all double-bit errors and all errors within a nibble.

The MPC8560 can invoke a level of system power management by invoking on-the-fly the CKE SDRAM signal to put the memory into a low-power sleep mode.

2.6 Interrupt Controller

The programmable interrupt controller 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 MPC8560 interrupt controller unit supports its single processor and provides for 12 external interrupts (with fully nested interrupt delivery), four message interrupts, internal-logic driven interrupts, and four 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 communication is supported through the external interrupt and soft reset signals of previous processors that implement the PowerPC architecture.

2.7 I^2C Controller

The inter-IC (IIC or I^2C) bus is a two-wire, bidirectional serial bus that provides a simple, efficient method of data exchange between devices. The The synchronous, multi-master bus of the I2C allows the MPC8560 to exchange data with other I^2C 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, multimaster bus of the I^2C allows the connection of additional devices to the bus for expansion and system development.

The I^2C controller is a true multimaster bus which 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 I^2C controller consists of a transmitter/receiver unit, a clocking unit, and a control unit. The I^2C unit supports general broadcast mode and on-chip filtering rejects spikes on the bus.

2.8 Boot Sequencer

The MPC8560 provides an boot sequencer that uses the I^2C interface to access an external serial ROM and loads the data into the MPC8560's configuration registers. The boot sequencer is enabled by a configuration pin sampled when the MPC8560's hardware reset signal is negated. If enabled, the boot sequencer holds The MPC8560's processor core in reset until the boot sequence is complete. If the boot sequencer is not enabled, the processor core is allowed to exit reset and fetch boot code from a port in default configuration.

2.9 Local Bus Controller

The MPC8560 local bus 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, FEPROM, ROM, and other devices such as 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 local bus interface unit provides a synchronous DRAM (SDRAM) machine that provides 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.

2.10 Triple Speed (10/100/1000) Ethernet Controllers

Each triple speed Ethernet controller (TSEC) incorporates a MAC that supports 10, 100, and 1000 Mbps Ethernet/802.3 networks with MII, GMII, RGMII, RTBI, and TBI physical interfaces. The TSEC includes address/data filtering, data insertion and extraction, 2-Kbyte FIFOs, and DMA functions.

The buffer descriptors are backwards compatible with the MPC8260 and MPC860T 10/100 programming models. Extended programming model features for the MPC8560 include four unicast MAC addresses, recognition of hash, broadcast, and multicast addresses, frame and extended buffer descriptors, and support for promiscuous mode.

The extended buffer descriptors support advanced features such as automatic data extraction and insertion.

Each MPC8560 TSEC supports programmable CRC generation and checking, RMON statistics, and jumbo frames up to 9.6 Kbytes.

2.11 Integrated DMA

The MPC8560 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 DMA 4-channel controller allows chaining (both extended and direct) through local memory-mapped chain descriptors. Transfers can be scatter gathered and unaligned. 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 or from external DMA request signals associated with each channel. There is also support for externally controlled transfer using $\overline{\text{DMA\_DREQ}}$, $\overline{\text{DMA\_DACK}}$, and $\overline{\text{DMA\_DDONE}}$.

DMA descriptors encompass a rich set of attributes that allow DMA transfers to bypass internal address translation and supply external addresses and attributes directory to any of I/O ports. Local attributes such as snoop and L2-write 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).

A packet descriptor mode allows the DMA controller to move data packets that are described by the same descriptors as those used by the triple speed Ethernet controller. This enables tight coupling between the Ethernet controller and the DMA controller and allows the DMA controller to efficiently forward incoming packets to a target.

The MPC8560 DMA can operate in stride mode so that a chunk of data can be transferred.

2.12 PCI Controller

The MPC8560 64-bit PCI controller is compatible with the *PCI Local Bus Specification, Revision 2.2* and the *PCI-X Addendum, Revision 1.0*. The PCI interface can function either as a host or agent bridge interface, and 64-bit dual address cycles are supported. The interface also supports devices with a 32-bit wide data bus.

The 64-bit PCI interface of the MPC8560 functions as both a master and target device. As a master, the MPC8560 supports read and write operations to the PCI memory space, the PCI I/O space, and the PCI configuration space. The MPC8560 also supports PCI special-cycle and interrupt-acknowledge commands. As a target, the MPC8560 supports read and write operations to system memory as well as configuration accesses. Arbitration occurs on the chip through a two-level priority round-robin arbitration algorithm. PCI reads are prefetched in memory. Coherency is selectable and hardware enforced.

PCI-X functionality includes split transaction support for one split.

2.13 RapidIO Controller

The RapidIO interconnect unit on the MPC8560 is based on the Rev. 1.1 “RapidIO Interconnect Specification.” 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 MPC8560 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, third party intervention, packet stomping, 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 1 Gbyte/s for an 8-bit port. Packet pacing

occurs at the physical layer. ATOMIC operations include increment, decrement, set, and clear.

The MPC8560 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.

2.14 Boot Sequencer

The MPC8560 provides a boot sequencer that uses the I^2C interface to access an external serial EEPROM and loads the data into the MPC8560 configuration registers. The boot sequencer is enabled by a configuration pin sampled during the reset sequence. If enabled, the boot sequencer completes the boot sequence before the core fetches its reset vector. A flexible programming format for the EEPROM content allows programming of any memory-mapped register or peripheral on or connected to the MPC8560. If the boot sequencer is not enabled, the processor core can exit reset and fetch boot code from a port in its default configurations. Three pins are used to determine the default boot source.

2.15 Address Map

The MPC8560 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 MPC8560 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 unit (ATMU). Both inbound and outbound translation windows are provided. The ATMU allows the MPC8560 to be part of larger address maps such as the PCI 64-bit address environment and the RapidIO environment.

2.16 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 switch fabric 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.

2.17 Data Processing Overview

Protocol data units (PDUs) can navigate through the various MPC8560 I/O ports in three ways. In the first, data is processed by the MPC8560 CPM (as it is received and transmitted through the UTOPIAs, MIIs, and TDMs associated with the CPM unit) and the local bus. In the second method, data is received by any of the available I/O ports, is sent through the

on-chip switch fabric, and is finally transmitted on the target I/O port without the use of the ECM. Lastly, data can be routed from any I/O port to any other I/O port via the ECM.

2.18 Processing Between the CPM and Local Bus

In this case, the MPC8560 stores data in buffers which reside in SDRAM on the local bus. These buffers are each referenced by a buffer descriptor (BD), which may reside in one of two tables (Rx-receive and Tx-transmit) typically placed in DPRAM. The following is a general overview of how incoming data is processed by the CPM (refer to Figure 2-3).

1. When Rx data arrives on the I/O port, it is decoded: the PDU is delineated from the incoming data stream.
2. Next, data is converted from its serial form into a parallel form and then loaded into the Rx FIFO.
3. When the Rx FIFO is filled to a set threshold, the respective communication channel signals the CPM for service.
4. The CPM accesses the next available RxBD in the RxBD table (pointed to by a register in the channel's parameter RAM table). The BD defines the main memory location where the data is to be placed, as well as the length of this buffer. Data is then moved from the Rx FIFO to a temporary storage location by the CPM.
5. Data is finally moved from this temporary storage location to main memory through DMA transactions. The status and control bits of the BD are updated, and the BD is closed.
6. A CPU interrupt is instantiated to notify the core that a new packet has been received.

Because FIFOs are typically smaller than the incoming PDU, steps 1–3 may be repeated several times to store the entire packet. There may also be times when the incoming PDU is larger than the buffer length defined in the RxBD. In such cases, steps 4–5 may be repeated and several BDs may be opened and closed to store the entire PDU.

When the transmit portion of the communication channel is enabled, the CPM starts with the first BD in the TxBD table (pointed to by a register in the channel's parameter RAM table), polling the ready R bit of the BD to verify that the buffer is ready for transmission.

7. When the BD is marked ready, the data is moved from the main memory buffer to a temporary storage location through DMA transactions.
8. The CPM moves data from the temporary memory location to the Tx FIFO.
9. Data is taken from the Tx FIFO in its parallel form and serialized.
10. When the Tx FIFO is emptied to a set threshold, the communication channel signals the CPM for more data in order to maintain throughput.
11. The serialized data is lastly encoded and transmitted on the I/O port.

Again, because the buffer pointed to by the Tx BD is typically larger than the Tx FIFO, the communication channel may make several iterations of steps 7–10 to load and transmit the entire PDU.

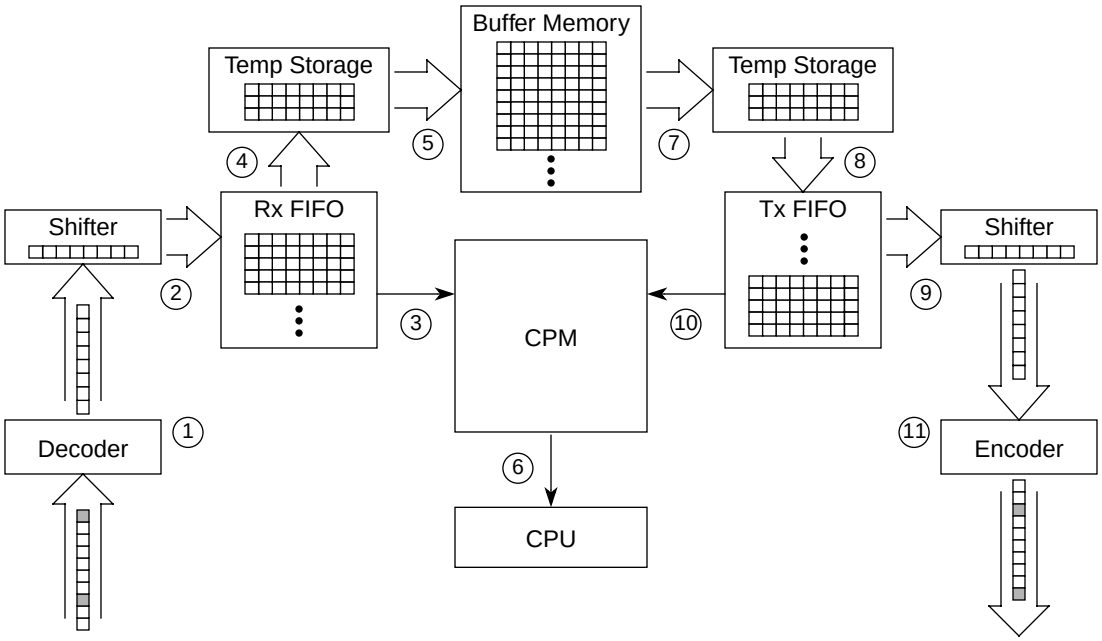

Figure 2-3. Data Processing Within the CPM

2.19 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 MPC8560. 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 ATMU processes transactions over the on-chip fabric (refer to Figure 2-4).

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.

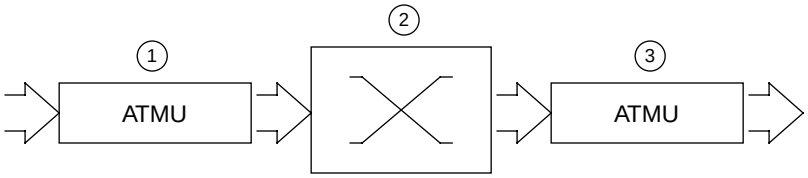

Figure 2-4. Processing Transactions Across the On-Chip Fabric

Figure 2-5 shows an example of how data is processed across the on-chip fabric. In this example the TSEC interacting with GMII I/O represents the module at the origination port of the fabric, and the destination buffer descriptors and data buffers are located in main memory on the PCI bus.

1. When Rx data arrives on the I/O port, it is decoded: the PDU is delineated from the incoming data stream. Next, data is converted from its byte form into a 64-bit word form and then loaded into a Rx FIFO.
2. The next available BD in the Rx BD table is accessed. It defines among other things the main memory destination location for the data. The ATMU at the TSEC fabric port takes the address referenced in the BD and generates a fabric port destination ID and a new local device address.
3. Data is then moved from the temporary storage location across the on-chip fabric to the destination port (PCI port).
4. The destination port's ATMU takes the local device address generated for this data and translates it to reflect the appropriate memory map on the PCI bus.
5. The data is then moved to the memory location defined in the BD, and the BD is updated and closed.

When the transmit portion of the communication channel is enabled, the first BD in the Tx BD table is polled to verify that the buffer is ready for transmission.

6. When the BD is marked ready, the ATMU at the PCI port takes the address referenced in the BD and generates a fabric port destination ID and a new local device address.
7. The data is moved from the data buffer across the on-chip fabric to the destination port (TSEC) defined by the ATMU at the PCI port.
8. The receiving ATMU at the destination port takes the local device address generated for this data and translates it to reflect the appropriate memory map on the destination port while moving the data into a Tx FIFO.
9. Data is converted from its 64-bit word form to the byte format needed for transmission and is lastly encoded and transmitted on the GMII I/O port.

It should be noted that BDs retrieved by the TSEC ATMU from main memory on the PCI bus go through the same translation process as data (similar to steps 6 and 7) as they are pulled across the on-chip fabric. It should also be noted that a single PDU may require more

than one BD for reference, in which case steps 2–5 may be repeated for receive transactions and steps 6–9 may be repeated for transmit transactions to completely move the entire PDU.

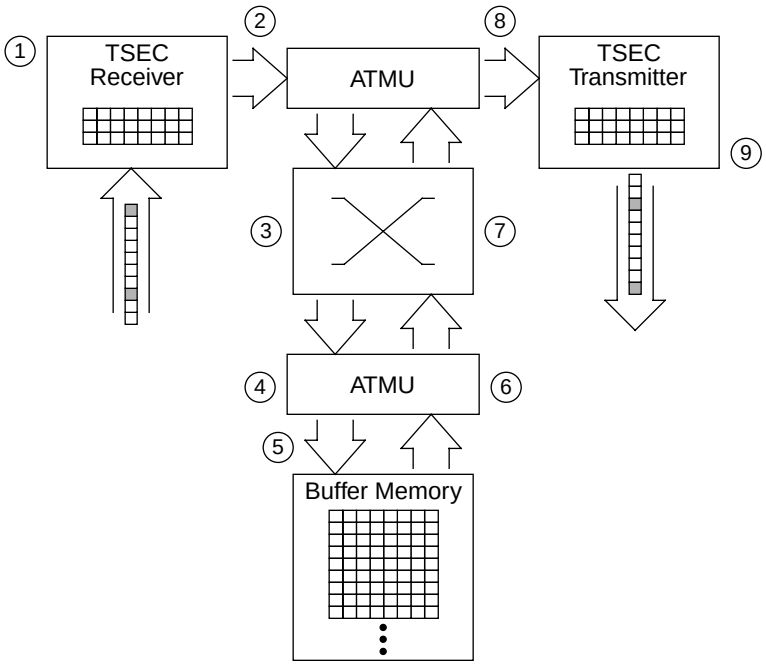

Figure 2-5. Example of Data Processing Across the On-Chip Fabric

2.20 Data Processing with the e500 Coherency Module

Processing via the ECM is similar to processing between the CPM and local bus or 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 (except in the case where the CPM is directly accessing the local bus). However, simply because the ECM is used 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 recognized 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 others (between the CPM and the local bus, and across the on-chip fabric) are non-coherent.

While transactions between the CPM and local bus are considered non-coherent because the CPM typically interfaces directly to the local bus (where its buffers are stored), CPM transactions can be made coherent. ATM transactions on a per-connection and direction basis can be set as coherent by setting the necessary bits in the receive and transmit connection tables. Coherency of MCC transactions per logical channel is determined by bits set in TSTATE. FCC and SCC transactions per physical channel and direction can be programmed as coherent by asserting the appropriate bits in the FCC and SCC functional code registers, respectively.

Part III MPC8560 Application Examples

The following section provides block diagrams of different MPC8560 applications. The MPC8560 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.

3.1 Device Configurations

There are three main system configurations for the MPC8560:

- Single-processor system
- Multiprocessor system
- High-performance system

3.1.1 Single-Processor System

In this system configuration, shown in Figure 3-6, the MPC8560 core uses the 64-bit DDR SDRAM bus to store data. The 32-bit local bus data is needed to store connection tables for many active ATM connections. The local bus may also be used to store data that does not need to be heavily processed by the e500 core. The CP can store large data frames in the local memory without interfering with the operation of the e500 core.

Figure 3-6 shows a basic system configuration.

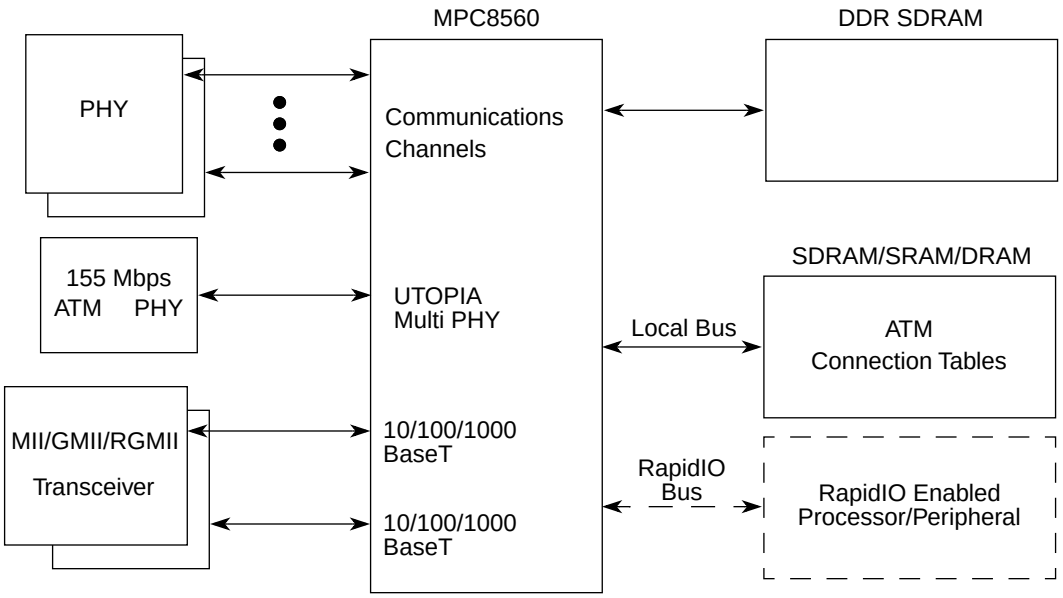

Figure 3-6. Basic System Configuration

3.1.2 Multiprocessor System

Figure 3-7 shows a multiprocessor system configuration.

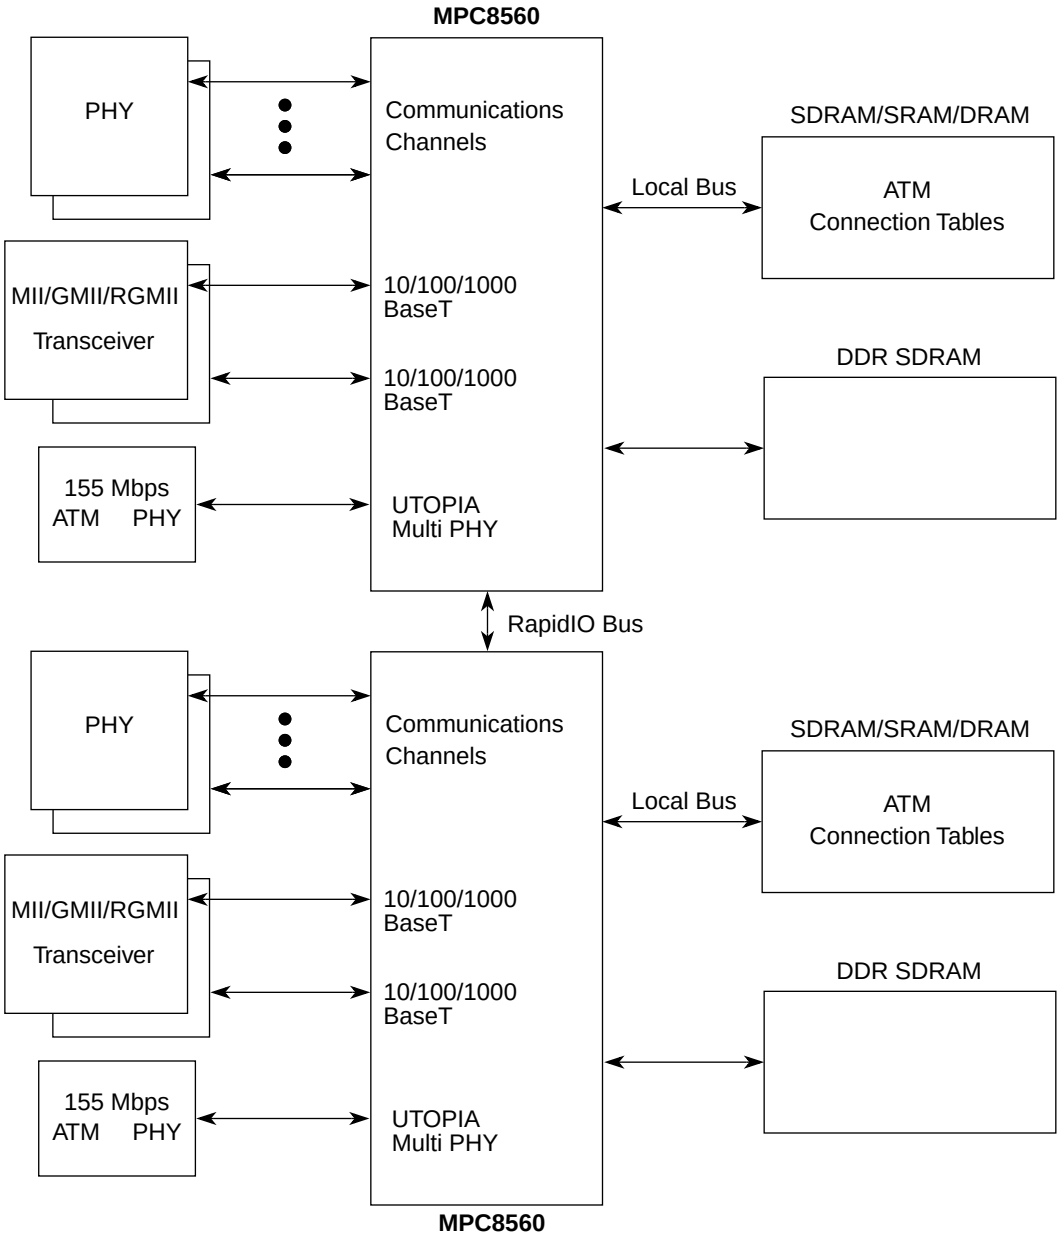

Figure 3-7. High-Performance Communications

This system enhances the serial throughput by connecting one MPC8560 to another MPC8560 with the RapidIO interface. The core in one of the MPC8560 devices can easily access the data stored in the DDR SDRAM memory of the other MPC8560. For performance reasons, the connection table information stored in the local bus memory for one MPC8560 should be copied into the local bus memory of the second. A system of this type can support 512 64-Kbps channels.

3.1.3 High-Performance System

Figure 3-8 shows a configuration with a high-performance system.

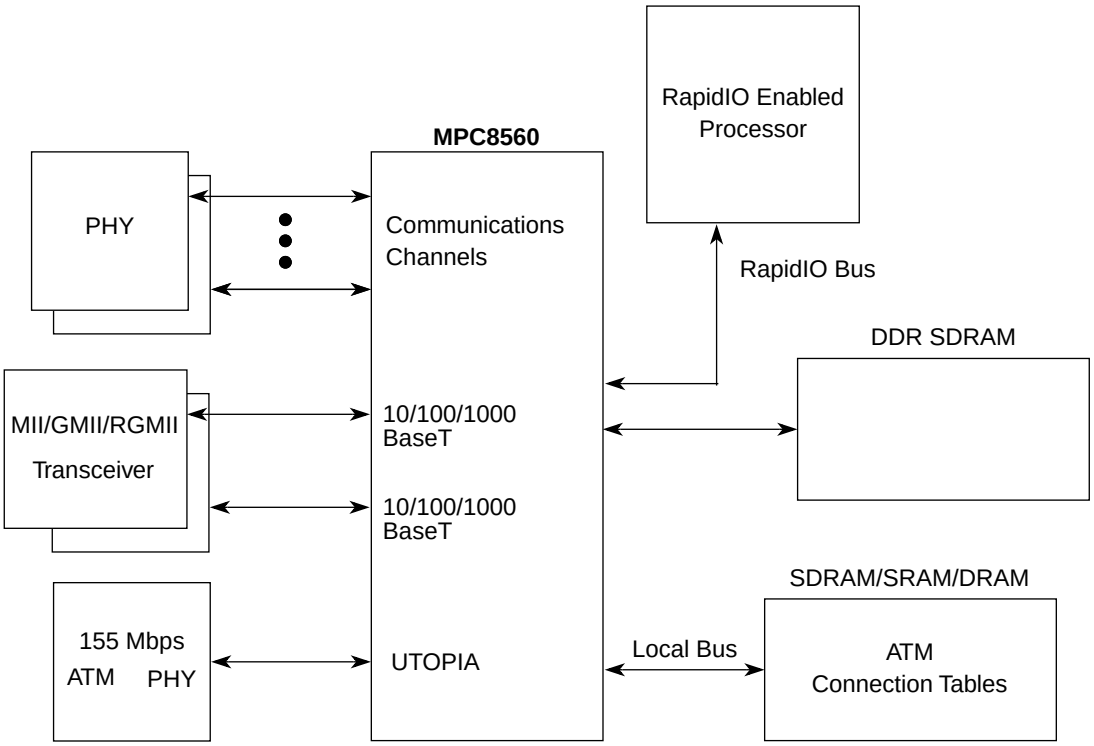

Figure 3-8. High-Performance System Microprocessor Configuration

In this system, an external high-performance microprocessor is connected to the MPC8560 via the RapidIO bus to share the processing load of the e500 core between the MPC8560 and the external processor to increase higher layer processing.

3.2 Examples of Communications Systems

Following are some examples of communications systems:

- Remote access server
- Regional office router
- LAN-to-WAN bridge router
- Cellular base station
- 3G wireless base station
- Telecom switch controller
- SONET transmission controller
- Frame relay card
- ATM protocol converter

3.2.1 Remote Access Server

See Figure 3-9 for a remote access server configuration.

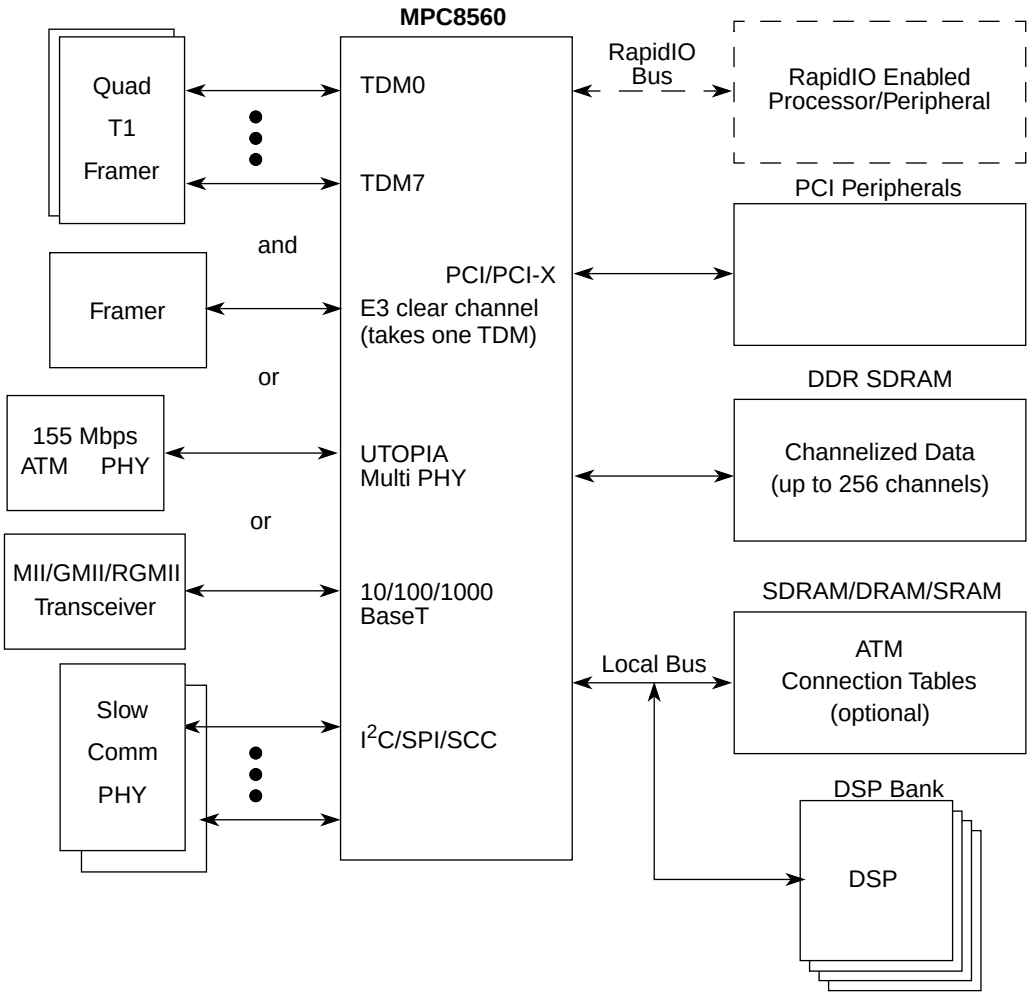

Figure 3-9. Remote Access Server Configuration

In this application, eight TDM ports are connected to external framers. In the MPC8560, each group of four ports support up to 128 channels. One TDM interface can support 32–128 channels. The MPC8560 receives and transmits data in transparent or HDLC mode, and stores or retrieves the channelized data from memory. The data can be stored either in the DDR SDRAM memory or in memory residing on the local bus.

The main trunk can be configured as one of the following:

- 155 Mbps full-duplex ATM, using the UTOPIA interface
- 10/100/1000BaseT Ethernet with MII/GMII/RGMII interface
- High-speed serial channel (up to 45 Mbps)

In ATM mode, there may be a need to store connection tables in external memory on the local bus (if more than 128 active connections are needed). The need for local bus memory depends on the total throughput of the system. The MPC8560 supports automatic (without software intervention) cross connect between ATM and MCC, routing ATM AAL1 frames to MCC slots.

The local bus can be used as an interface to a bank of DSPs that can perform analog modem signal modulation. Data to and from the DSPs can be transferred through the MPC8560 DMA controller.

The MPC8560 local bus memory controller supports pipeline SDRAM devices for efficient burst transfers. A separate DDR controller and bus is available on the MPC8560 to support DDR SDRAM.

3.2.2 Regional Office Router

Figure 3-10 shows a regional office router configuration.

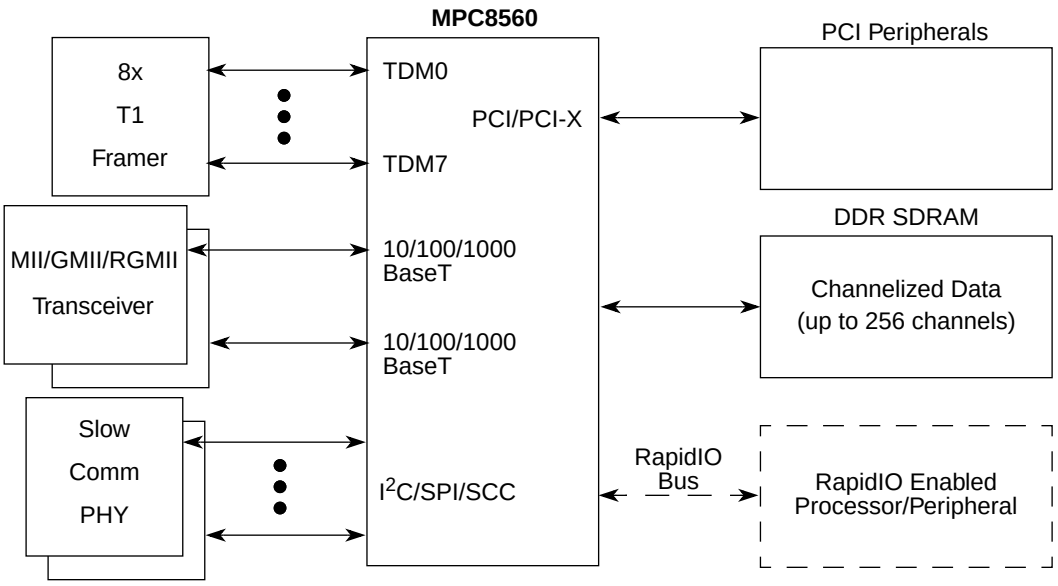

Figure 3-10. Regional Office Router Configuration

In this application, the MPC8560 is connected to up to eight TDM interfaces with up to 256 channels. Each TDM port supports 32–128 channels. If 256 channels are needed, each TDM port can be configured to support 32 channels. In this application there are two MII/GMII/RGMII ports for 10/100/1000BaseT LAN connections.

In all the examples, the SCC ports can be used for management.

3.2.3 LAN-to-WAN Bridge Router

Figure 3-11 shows a LAN-to-WAN router configuration, which is similar to the previous example.

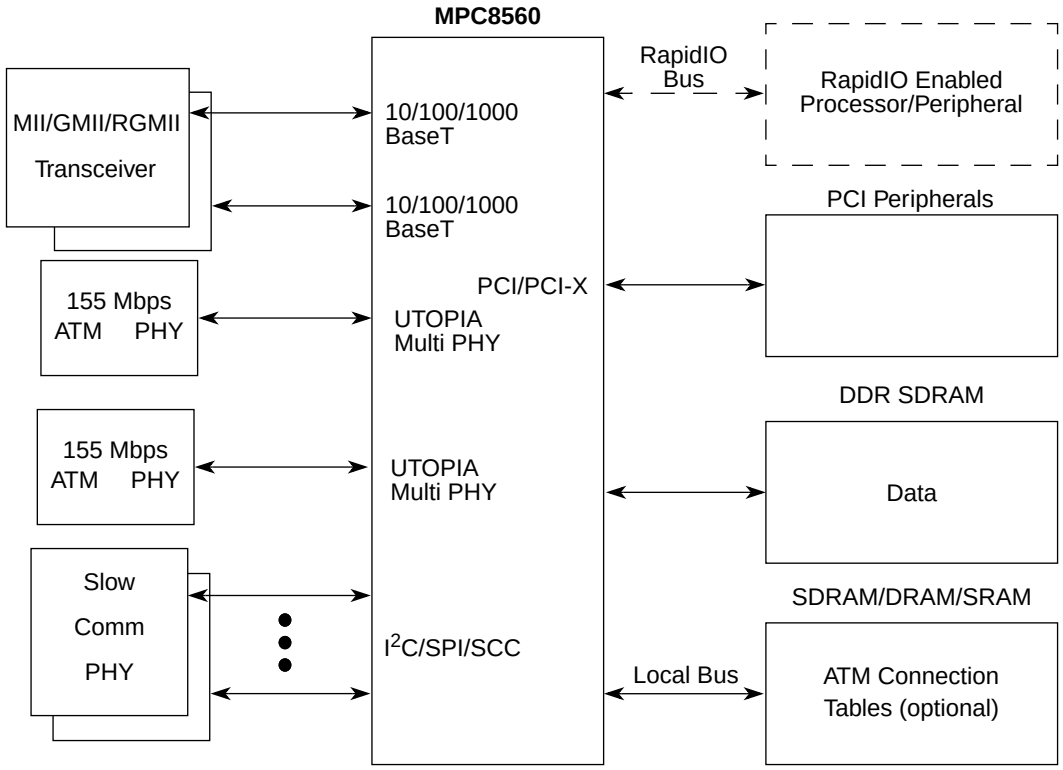

Figure 3-11. LAN-to-WAN Bridge Router Configuration

3.2.4 Cellular Base Station

Figure 3-12 shows a cellular base station configuration.

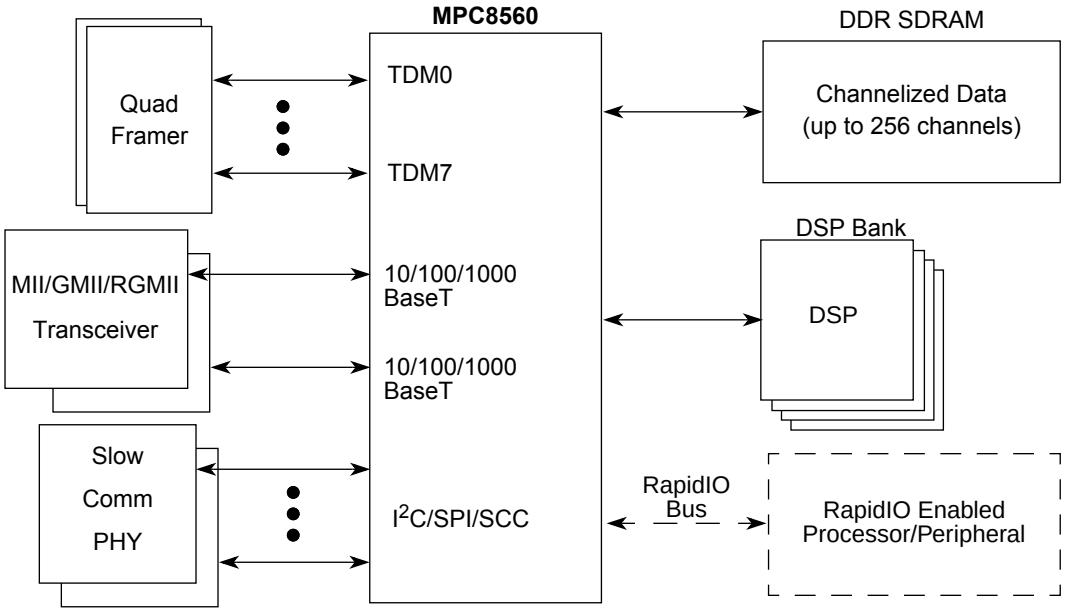

Figure 3-12. Cellular Base Station Configuration

Here the MPC8560 channelizes eight E1s (up to 256, 64-Kbps channels). The local bus can control a bank of DSPs. Data to and from the DSPs can be transferred through the local bus to the host port of the DSPs through the integrated DMA controller. The slower communications ports (SCCs, I^2C, SPI) can be used for management and debug functions.

3.2.5 3G Wireless Base Station

Figure 3-13 shows a 3G wireless base station configuration.

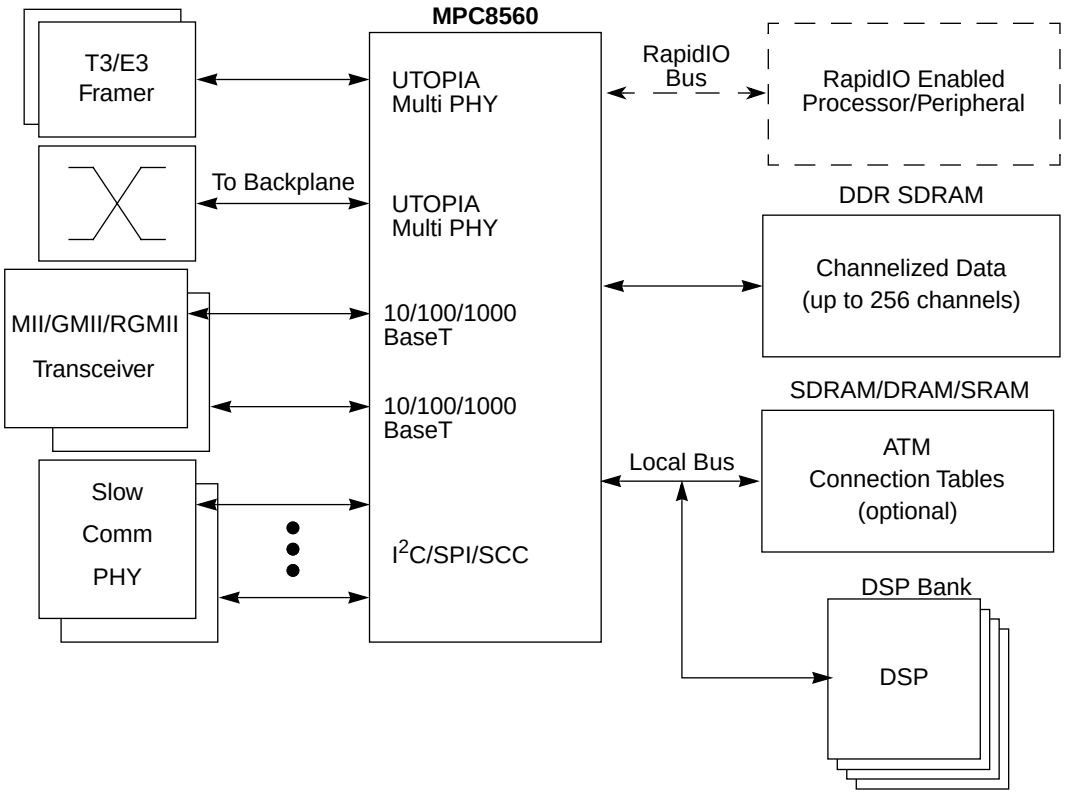

Figure 3-13. 3G Wireless Base Station Configuration

Here the MPC8560 uses two E3/T3s for ATM connections or alternatively two high-bit rate HDLC connections. The local bus can control a bank of DSPs as well as store ATM connection tables. Data to and from the DSPs can be transferred through the local bus to the host port of the DSPs through the integrated DMA controller. The slower communications ports (SCCs, I^2C, SPI) can be used for management and debug functions.

3.2.6 Telecommunications Switch Controller

Figure 3-14 shows a telecommunications switch controller configuration.

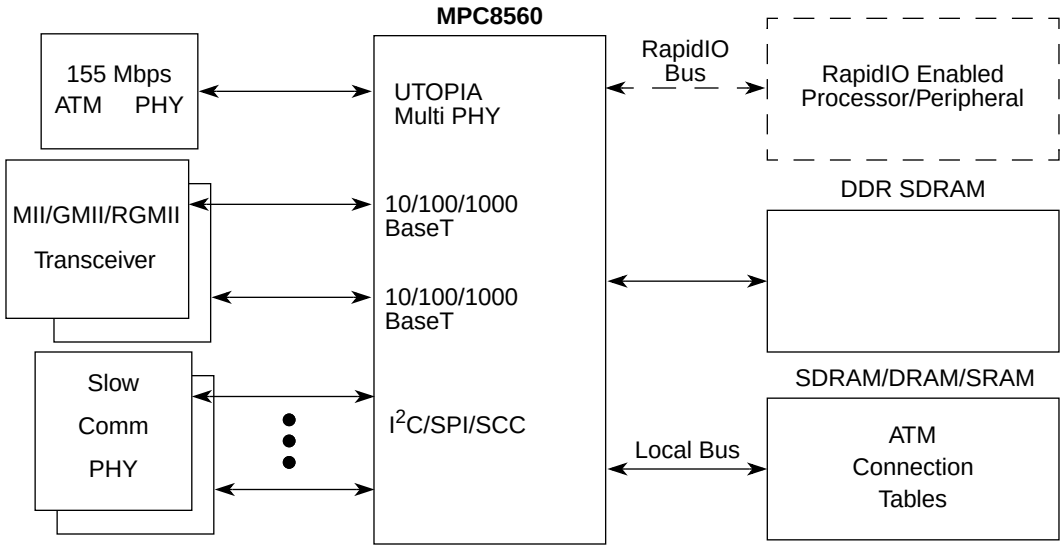

Figure 3-14. Telecommunications Switch Controller Configuration

The MPC8560 CPM supports a total aggregate throughput of 1 Gbps at 333 MHz. This includes one full-duplex, 155-Mbps ATM channel. For more information, reference the CPM section of the *MPC8560 User's Manual*.

3.2.7 SONET Transmission Controller

Figure 3-15 shows a SONET transmission controller configuration.

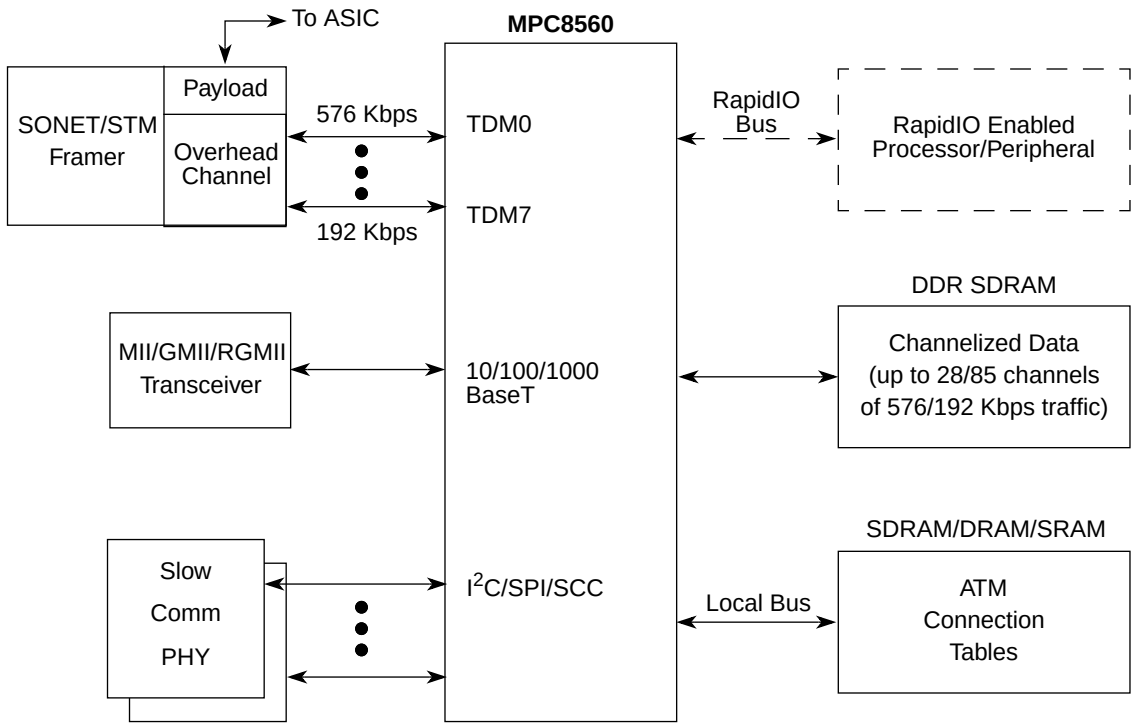

Figure 3-15. SONET Transmission Controller Configuration

In this application, the MPC8560 implements super channeling with the multichannel controller (MCC). Nine 64-Kbps channels are aggregated to form a 576-Kbps channel. The MPC8560 at 333 MHz can support up to 28 576-Kbps superchannels. The MPC8560 also supports subchanneling (under 64 Kbps) with its MCC.

3.2.8 Frame Relay Card

Figure 3-16 shows a frame relay card configuration.

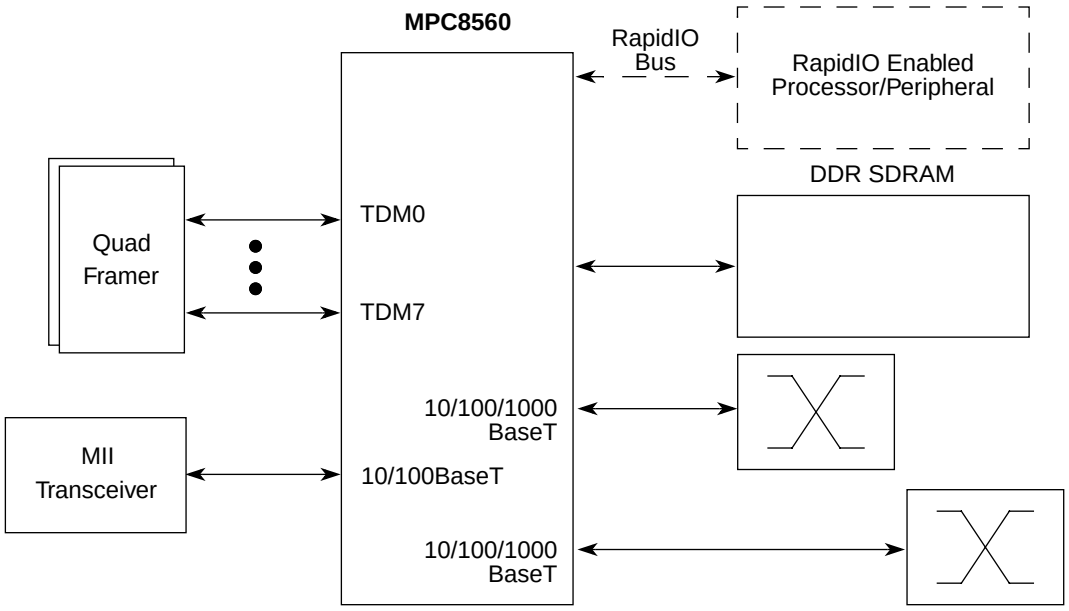

Figure 3-16. Frame Relay Card Configuration

3.2.9 ATM Protocol Converter

Figure 3-17 shows an ATM protocol converter.

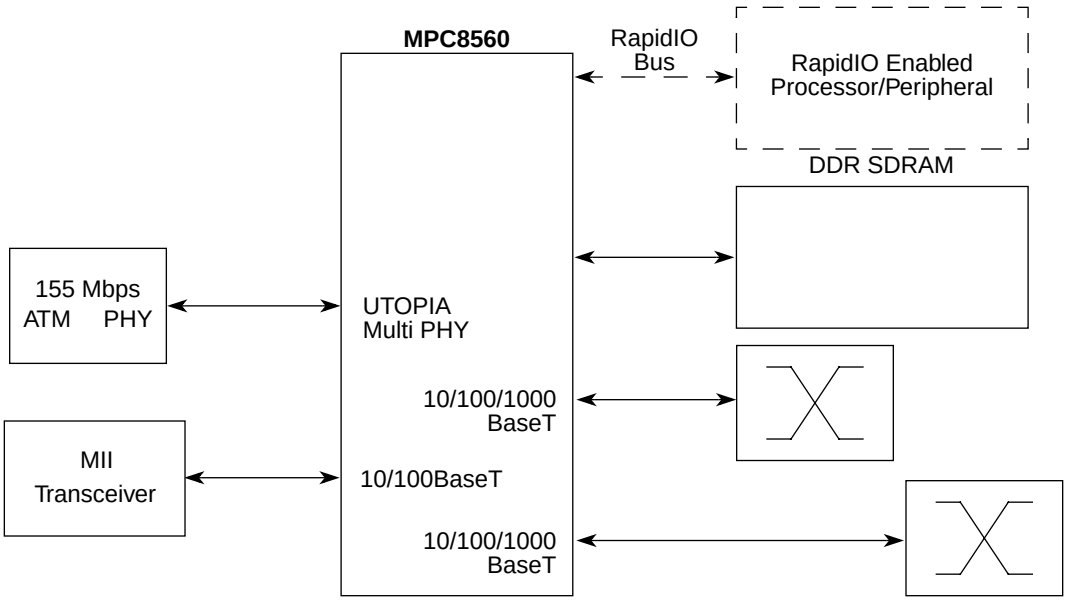

Figure 3-17. ATM Protocol Converter Configuration

In this configuration the MPC8560 can convert traffic from ATM to Ethernet and from Ethernet to ATM. The MPC8560 is also connected to a redundant gigabit Ethernet switch fabric backplane via the two TSECs.

Part IV Compatibility Issues

This section describes some software and hardware compatibility issues.

4.1 Software

The MPC8560 CPM features are similar to those in the previous generation MPC8260. The code ports easily from previous devices to the MPC8560, except for new protocols. During the definition of this device, an effort was made to maintain compatibility wherever possible.

Note that the MPC8560 initialization code requires changes from the MPC8260 initialization code (Motorola will provide the reference code.)

4.2 MPC8560 Hardware

As the MPC8560 family migrates to smaller geometries, the core voltage will reduce from 1.2 V to low voltages. A programmable voltage regulator is recommended for future compatibility. See the MPC8560 Hardware Specifications for the electrical requirements and the AC and DC characteristics.

4.3 Differences Between MPC8560 and MPC8260

While the MPC8560 adds considerably to the functionality offered by the MPC8260, the following MPC8260 features are not included on the MPC8560:

- 60x bus interface
- CPM IDMA support and FlyBy DMA. (An IDMA controller outside the CPM is replacing this.)
- IEEE 802.3/Ethernet support on the SCCs
- SMC functional blocks in the CPM

4.4 Communications Protocol Table

Table 4-1 summarizes available protocols for each communications port.

Table 4-1. MPC8560 Protocols

Protocol	Port				
	TSEC	FCC	SCC	MCC	TC Layer
ATM (UTOPIA)		√			
ATM (Serial)					√
1000BaseT	√				
100BaseT	√	√			
10BaseT	√	√			
HDLC		√	√	√	
HDLC_BUS			√		
Transparent		√	√	√	
UART			√		
Multichannel				√	

4.5 MPC8560 Configurations

The MPC8560 offers flexibility in configuring the device for specific applications. The functions mentioned in the above sections are all available in the device, but not all of them can be used at the same time. This does not imply that the device is not fully activated in any given implementation. The CPM architecture has the advantage of using common hardware resources for many different protocols, and applications. Two factors limit the functionality in any given system: pinout and performance.

4.5.1 Pin Configurations

To maximize the efficiency of device pins, some pins have multiple functions. In some cases choosing a function may preclude the use of another function.

4.5.2 Communications Performance

The CPM is designed to handle an aggregate of 1 Gbps on the communications channels running at 333 MHz. Performance depends on a number of factors:

- Channel rate versus CPM clock frequency for adequate polling of communications channels for service
- Channel rate and protocol versus CPM clock frequency for CP protocol handling
- Channel rate and protocol versus bus bandwidth
- Channel rate and protocol versus system core clock for adequate protocol handling

The second item above is addressed in this section—the CP's ability to handle high bit-rate protocols. Slow bit-rate protocols do not significantly affect those numbers.

Table 4-2 shows the peak CPM performance of various protocols under the assumption that only one of those protocols is running at a given time. The ATM numbers shown also assume that the local bus is used exclusively by the CPM and enough bandwidth on the DDR memory system is available for the CPM (implying that other resources like the PCI-X controller, TSECs, RapidIO interconnect, DMA controller, and CPU do not all operate at their maximum performance). The frequency specified is the minimum CPM frequency necessary to run the mentioned protocols concurrently in full-duplex.

Table 4-2. MPC8560 Serial Performance

<table>
<tr><th colspan="3" rowspan="2">Protocol</th><th colspan="3">Frame Size</th></tr>
<tr><th>1024 Bytes</th><th>128 Bytes</th><th>64 Bytes</th></tr>
<tr><td>Ethernet</td><td colspan="2">FCC: 100BaseT</td><td colspan="3">3 x 100BaseT Full Duplex = 600 Mbps</td></tr>
<tr><td>HDLC</td><td colspan="2">MCC: HDLC</td><td colspan="3">256 Full Duplex Channels at 64 Kbps = 16.7 Mbps</td></tr>
<tr><td rowspan="4">ATM</td><td rowspan="2">FCC: AAL5</td><td>No Bus Limitation</td><td colspan="2" rowspan="2">≥ 1000 Mbps Aggregated</td><td rowspan="2">≥ 900 Mbps Aggregated</td></tr>
<tr><td>Connection Tables on Local Bus</td></tr>
<tr><td rowspan="2">FCC: AAL0</td><td>No Bus Limitation</td><td colspan="3" rowspan="2">≥ 1000 Mbps Aggregated</td></tr>
<tr><td>Connection Tables on Local Bus</td></tr>
<tr><td colspan="2">FCC: AAL2 CPS</td><td colspan="3">33–242 Mbps Depending on PDU Size</td></tr>
</table>

These performance estimates assume the CPM is operating at 333 MHz, and that data is stored in SDRAM on the local bus operating at 166 MHz.

HOW TO REACH US:

USA/EUROPE/LOCATIONS NOT LISTED:

Motorola Literature Distribution
P.O. Box 5405, Denver, Colorado 80217
1-303-675-2140 or 1-800-441-2447

JAPAN:

Motorola Japan Ltd.
SPS, Technical Information Center
3-20-1, Minami-Azabu Minato-ku
Tokyo 106-8573 Japan
81-3-3440-3569

ASIA/PACIFIC:

Motorola Semiconductors H.K. Ltd.
Silicon Harbour Centre, 2 Dai King Street
Tai Po Industrial Estate, Tai Po, N.T., Hong Kong
852-26668334

TECHNICAL INFORMATION CENTER:

1-800-521-6274

HOME PAGE:

http://www.motorola.com/semiconductors

DOCUMENT COMMENTS:

FAX (512) 933-2625
Attn: RISC Applications Engineering

MOTOROLA

MPC8560TS/D