M X T 3 0 1 0

High-Performance Cell Processor

Typical Applications

Segmentation and Reassembly (SAR)

- ATM/packet/TDM interworking
- Service adaptation to cell-based backplanes at up to 800Mbps
- Combined SARing and cell-switching in low port-density applications

ATM port processing

- Traffic shaping and policing

WAN access "common card" services using the MXT3020 coprocessor

- AAL1, AAL2, AAL5
- Inverse Multiplexing for ATM (IMA), Cell Relay, Circuit Emulation Service, T1/E1 UNI

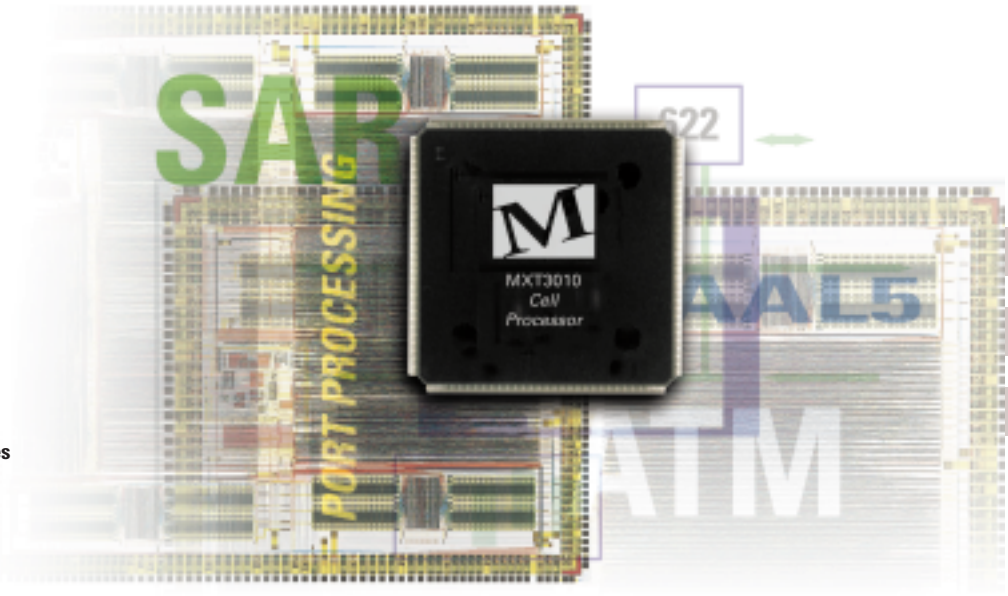

Distinctive Characteristics

Leading-edge performance

- Data rates up to 800Mbps
- Bi-directional OC-3 or unidirectional OC-12 at full line rates with minimum packet size
- Zero-time context switching

Complete programmability

- Adapts to new or changing standards via firmware upgrades
- Supports custom features

Concurrent support of CBR, VBR-nrt, VBR-rt, UBR, and ABR traffic classes

Full Suite of Development Tools

The MXT3010 is the industry's highest-performance programmable cell processing engine. It offers the throughput normally associated with hardwired designs while delivering all of the benefits of a programmable solution. Capable of supporting data rates up to 800 Mbps, the MXT3010 can perform data path functions such as segmentation/reassembly or traffic shaping for ATM interfaces running at up to OC-12 speeds. With its efficient RISC processor, high-speed context switching, and flexible cell scheduling system, the MXT3010 can easily integrate all current and emerging ATM Forum traffic classes, including ABR, in a single application.

Developers can license prepackaged software products that run on the MXT3010 or may write their own custom software. MXT3010 software products include:

CellMaker®-622: AAL5 SAR firmware supporting 1 x OC-12 or 4 x OC-3 configurations with integrated ABR traffic management.

CellMaker®-155: AAL5 SAR firmware supporting 1 x OC-3 or lower data rates, with integrated ABR traffic management.

CircuitMaker™: firmware for ATM/TDM/packet interworking using the MXT3010 Cell Processor and MXT3020 Circuit Coprocessor.

LinkMaker™: firmware for cost-effective Inverse Multiplexing for ATM.

MAKER COMMUNICATIONS, INC.

Segmentation and Reassembly

High Performance

The MXT3010 contains all of the necessary hardware support mechanisms to construct a high-performance SAR. CRC-32 and CRC-10 calculations are performed in hardware during data transfers. Independent DMA engines for each interface operate concurrently to achieve high data throughput by dispatching and pipelining operations. The cell scheduling system works with a combination of software algorithms and hardware-based bandwidth-reservation and request-resolution mechanisms to support multiple classes of service across up to 16,000 virtual connections.

Versatile

The MXT3010 adapts easily to both OC-3 and OC-12 rate SAR applications. An OC-3 SAR can be built using a single MXT3010 and consolidating user data and host control traffic across the 1.6 Gbps Port1 interface in a shared memory subsystem. At OC-12, higher data rates require specialization of the Port1 and Port2 interfaces as described in the caption to the left. In addition, two MXT3010s are required, each supporting half-duplex operation, to provide full-duplex OC-12 capability.

The MXT3010 is not limited to operation at only these data rates. With cell scheduling and traffic shaping controlled by firmware, virtually any data rate up to 800 Mbps can be supported.

Customizable

Because a RISC core controls the MXT3010, custom cell sizes or formats can be accommodated, as can proprietary traffic shaping algorithms and adaptation layers. Excess CPU bandwidth, available in OC-3 applications, can be used for enhanced statistics, protocol processing, or other host processor functions.

OC-3/OC-12 SAR.

The MXT3010 provides exceptional traffic management in the form of per-VC scheduling, the ability to support a variety of custom traffic shaping algorithms, and integrated RM cell processing for ABR. The diagram below illustrates that at OC-12 speeds, it is advantageous to separate data traffic from host control traffic. This way the high performance packet memory can be optimized for the burst sizes and rates common to segmentation and reassembly, without incurring overhead due to small, unpredictable command and response messages.

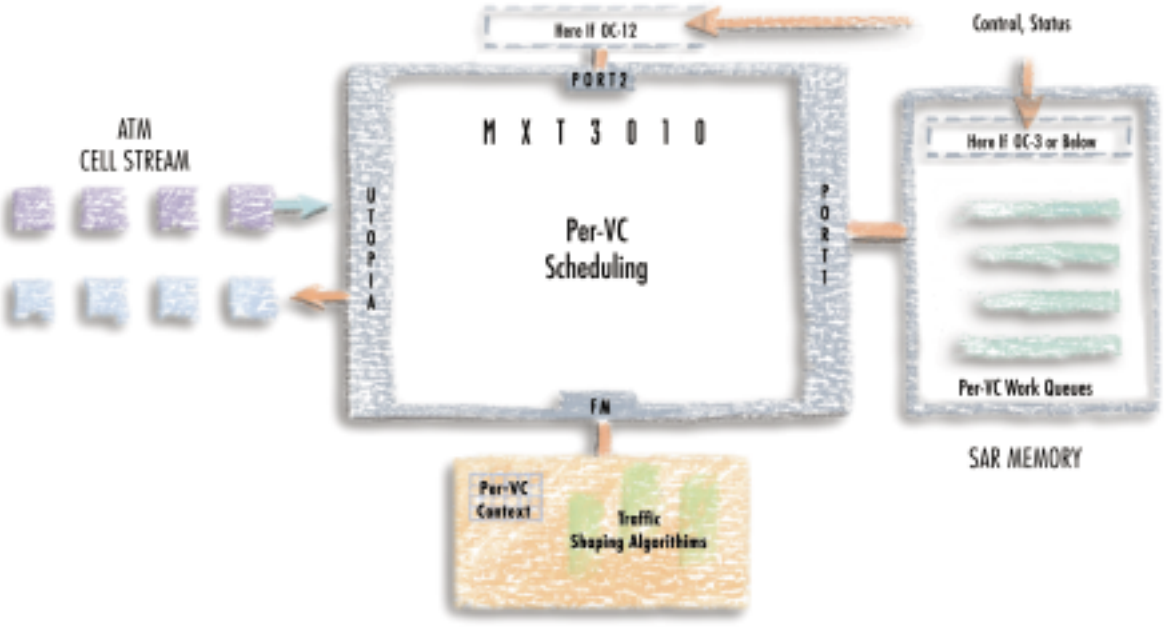

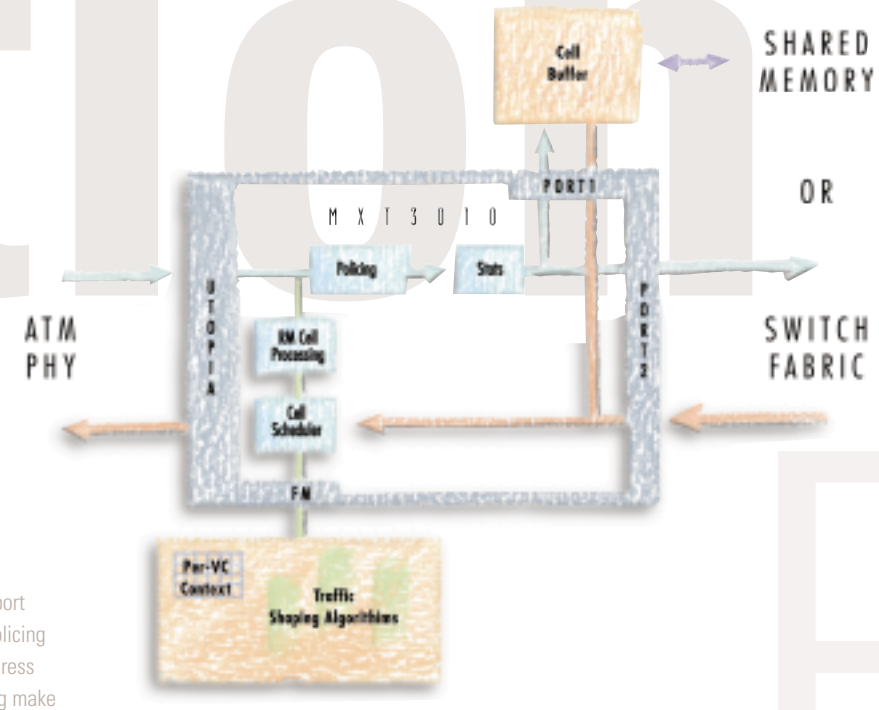

Single-point flow control.

The ability to combine ingress port processing functions such as policing and statistics gathering with egress functions such as traffic shaping make the MXT3010 an ideal device to perform flow control. The MXT3010 can support a variety of traffic shaping algorithms and can perform per-VC scheduling for up to 16,000 VCs.

Traffic Shaping

Traffic Shaping relies on two fundamental mechanisms: per-virtual circuit queuing and a Cell Scheduling System (CSS).

The MXT3010 simplifies the implementation of independent queues for tens of thousands of virtual circuits with a powerful context maintenance model. Separate channel descriptors for each VC provide the parameters necessary to track queue location and depth. The context switching architecture of the MXT3010 ensures that the channel descriptors can be loaded, manipulated and saved in the short cell interval time of OC-12 rates.

The chip performs scheduling through a combination of firmware algorithms running on the SWAN core and a hardware bandwidth-reservation mechanism. By creating service routines that are subsequently linked to by specific VCs, multiple separate traffic shaping algorithms (leaky bucket VBR, ABR, etc.) can be co-resident in a port processing application. Each VC activates its associated service routine during cell processing to calculate the appropriate inter-cell interval for transmission bandwidth reservation. These calculations can be complex in nature since they are handled by firmware running on the SWAN core. Virtual Source/Virtual Destination ABR algorithms can be implemented on the MXT3010. Actual transmission bandwidth is reserved through calls dispatched by the SWAN to the Cell Scheduling System. Each request of the CSS identifies a target transmission time. The CSS evaluates the current bandwidth reservation state as maintained in a schedule table, and assigns the next appropriate opportunity to the requesting VC. This assignment is recorded in the CSS table, which ultimately drives the ordering of the VCs serviced for transmission.

Inverse Multiplexing for ATM (IMA)

IMA involves splitting and sending a high bandwidth cell stream across multiple lower bandwidth links, and then reestablishing the original cell ordering at the receiver. It does this by inserting a framing and control protocol between the PHY and ATM layers at both endpoints of the IMA connection. The IMA control protocol cells contain information about IMA group status, IMA link status, loopback test information, and link synchronization information.

An IMA engine needs to operate in the ATM cell stream at line rate. It must recognize IMA control cells, extract those cells from the cell stream, process the cells, and based on the cells' content, adjust the IMA behavior.

The MXT3010 provides an ideal platform for this engine. The programmability of the device insures that an IMA system can evolve with changing standards. In addition to the basic IMA function, the MXT3010 can provide advanced features such as traffic shaping using per-VC queues and congestion management mechanisms such as early or partial packet discard.

A complete Ethernet-to-ATM edge device can be built with the MXT3010 by incorporating AAL5 SAR functions in the same application as the IMA protocol. In this manner, Ethernet packets can be converted to ATM cells and transmitted over an IMA group using a single MXT3010. A significant cost advantage is gained by combining these functions into a single device.

Single Chip Solution.

The MXT3010 is unique in its ability to perform both segmentation and reassembly and IMA processing in a single integrated solution.

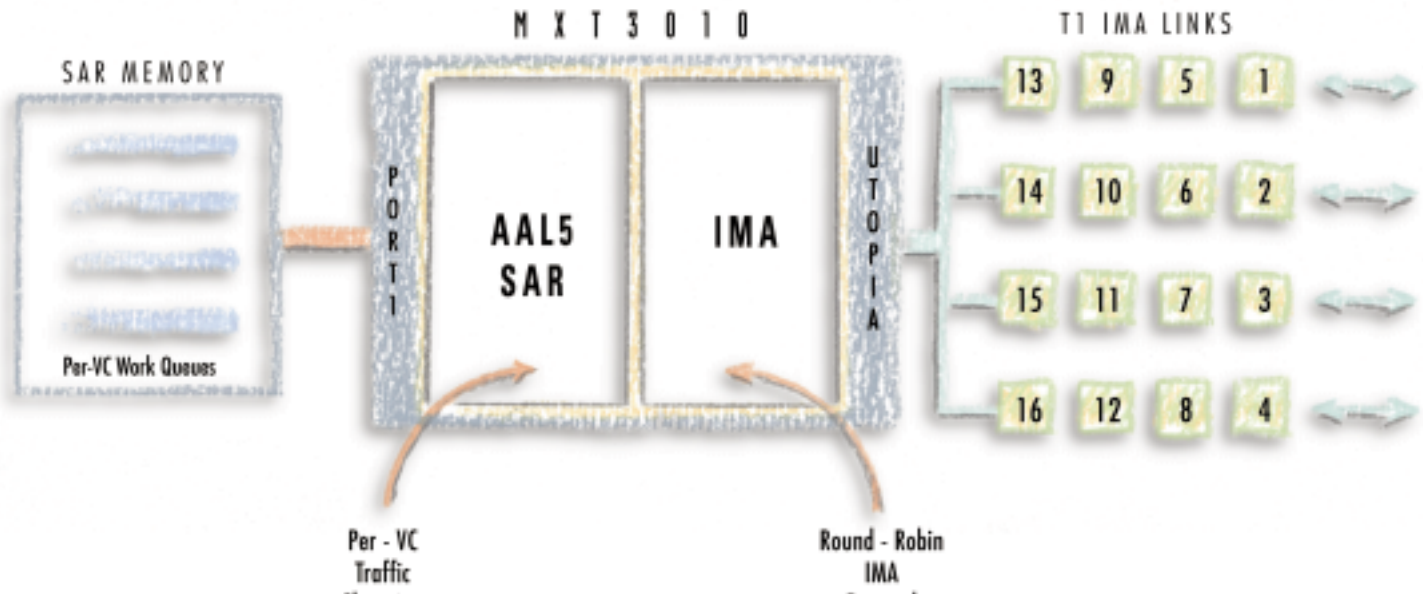

MXT3010 Features

SWAN™ Processor

- Pipelined, single-cycle operations
- Zero-time context switching
- 32-bit instruction set and 16-bit ALU
- Modulo arithmetic and built-in ALU branching
- Scoreboarded register set
- On-chip instruction cache
- Operation at 100, 80, or 66 MHz

On-Chip Hardware Agents

- Three DMA engines
- Cell Scheduling System
- Cell buffer RAM

Cell Transmission

- Per-VC scheduling
- 8 independent schedule tables for ports or Virtual Paths
- Software-controlled queuing, with a library of queuing functions

ATM Adaption Layer

- VPI/VCI reduction mechanism
- CRC-32 generation and checking
- CRC-10 generation and checking

ATM PHY Interface

- UTOPIA Level 2 Multi-PHY
- 8-bit bidirectional or 16-bit unidirectional
- Programmable HEC insertion
- Clocked independently from other system blocks

Theory of Operation

The MXT3010 delivers exceptional system performance through the use of parallel hardware agents, zero-time context switching, and an instruction set optimized for ATM cell manipulation and transmission. The chip is built around the SWAN™ (Soft-Wired ATM Network) processor, which is a specialized cell processing engine that supports data rates of up to 800 Mbps by combining the pipelined architecture of a RISC processor with the instruction set power of an ATM-specific CISC processor.

A number of hardware machines are integrated on-chip and operate in parallel to assist the SWAN. Three fully independent DMA engines work concurrently to sustain high throughput of data and control traffic through a high-speed multi-port internal cell buffer RAM. The scheduling of this traffic is managed continuously by an internal Cell Scheduling System (CSS).

The CSS includes a unique hardware scoreboard that can resolve scheduling conflicts for up to 16,000 active connections. This deterministic behavior is essential when maintaining independent traffic shaping parameters for each virtual connection in an OC-12 link. For more detail on the CSS, see the traffic shaping application example in this product brief.

External bus interfaces are available for each of the three DMA engines which access the internal cell buffer RAM. The UTOPIA port connects to an ATM network through a UTOPIA Level 2 multi-PHY interface. Port1 is a 32-bit DMA system interface which runs a lightweight, burst-mode protocol designed to deliver very high throughput while being easily adaptable to standard busses such as PCI or the i960. This port connects to packet memory in SARing applications or to an external cell buffer in port processing applications. Port2 is a 16-bit general-purpose interface, which supports both burst- and non-burst-mode operations. This port can be used for host control messages, allowing Port1 memory to be optimized for handling network traffic, or for connecting to MXT3020 Circuit Coprocessors in TDM applications. An inter-chip communications subsystem provides access to internal and external state information for application-specific use.

Ordering Information

MXT3010EP-A: 100 MHz MXT3010

MXT3010EP-A80: 80 MHz MXT3010

MXT3010EP-A66: 66 MHz MXT3010

MXT3010 Development Tools

Evaluation Cards: The MXT3025 Evaluation Board is available with a SONET OC-3 interface and four T1 interfaces. The MXT3016 Evaluation Board operates as a full-duplex OC-12 interface. Both cards are available with operational firmware.

CodeMaker™ Development Kit: This is a complete set of development tools including simulation environments, an assembler, a debugger, Verilog and Bus Functional Models.

Maker Communications, Inc.

73 Mount Wayte Avenue

Framingham, MA 01702

ph 508.628.0622

fx 508.628.0256

www.maker.com

Product Information:

info@maker.com

Summary Characteristics

Power Supply: 3.3V

I/O Voltage Levels: 3.3V- and 5V-compatible

Supply Current: 970 mA max @ 100 MHz and 3.3V

Maximum Power Dissipation: 2.1W @ 66 MHz, 2.6W @ 80 MHz, 3.2W @ 100 MHz

Typical Power Dissipation @ 100 MHz: 2.9W

Maximum Operating Junction Temperature: 110 °C

ESD Rating: 3KV

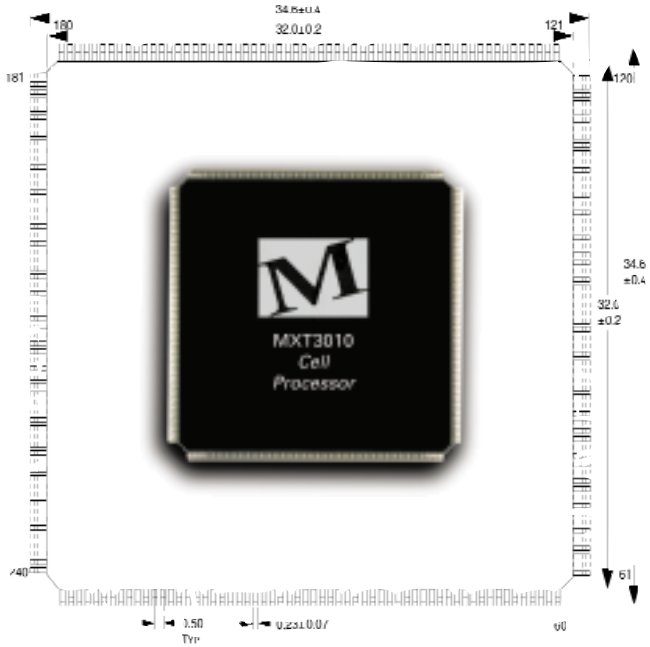

Physical Package		
Package Type	Body Size (mm)	Lead Pitch (mm)
PQFP 240	32.0 x 32.0 x 3.0	0.5

Moisture Sensitivity Rating	θjc (°C/W)	θja (°C/W)	
		*Still Air	*Air Flow (1m/sec)
JEDEC Level 3	2.0	20	14

* *These notes are for the packaged IC only. Actual results when soldered on to a printed circuit card will generally be lower and will vary depending on board stack-up and orientation. Refer to MXT3010 EP thermal test report.*