


RapidIO Architecture: Building the Next-Generation Networking Infrastructure



► Introduction

RapidIO[®], a high-performance, packet-switched bus technology, delivers the bandwidth, software independence, fault tolerance and low latency needed for the design of next-generation networking equipment in the communications market.

Digital hardware and embedded software designers in the communications industry face a challenge introduced by the increase in processor speeds and the bandwidth available to feed data to the processor. The critical bottleneck lies in the speed at which various components “inside the box” communicate with each other.

RapidIO, an open communications standard designed for chip-to-chip and board-to-board connections using a backplane, eliminates this bottleneck. Its architecture allows chip-to-chip and board-to-board transfer of data and control information within a chassis, far surpassing the bandwidth limits of legacy bus technologies, such as PCI, and significantly increasing transmission speeds.

RapidIO Specifications

The RapidIO Trade Association directs future development and promotion of the RapidIO interconnect as an open standard. The Association's Steering Committee oversees a Technical Working Group (TWG) and a Marketing Working Group (MWG).

Development, management and maintenance of specifications – an ongoing effort – is performed by technical subcommittees under the direction of the Steering Committee and TWG. These specifications provide for an extensive range of data widths and clock rates, ranging from 2.5 Gbits per link per direction for the 1x serial up to 32 Gbits per direction for 16-bit parallel with a 1 GHz clock. The most common implementations expected among the several defined in the RapidIO specifications include:

- 8-bit parallel with a 500 MHz clock (8 Gb/s per direction)
- 1x serial (2.5 Gb/s per direction)
- 4x serial (10 Gb/s per direction)

Parallel and Serial Implementations

The parallel versions of the RapidIO specification offer the highest bandwidth, lower overhead, lower latency and are intended for short transmission distances, whereas the serial configurations are intended for longer transmission distances and applications requiring the lowest pin counts and low power.

Specifications	Parallel	Serial
Bus/link Bandwidth	16, 32 Gb/s	1 to 10 Gb/s
Bus/link Full Duplex Bandwidth	32 to 64 Gb/s	2 to 20 Gb/s
Pin Count	40 to 76	4 to 16

Revision 1.1 of the parallel open standard specification is available to the public today. Release of the serial standard specification is scheduled in late 2002.

For more information about the RapidIO Trade Association and RapidIO specifications, please refer to the Association's website, www.rapidio.org.

RapidIO Applications

The primary application for RapidIO technology is the communication fabric inside networking equipment, both at the branch office for data transfer and control information, and in public-scale networking on the control side only. Initial target markets for RapidIO include networking and communications equipment, enterprise storage, and other high-performance embedded markets, while potential markets are the embedded military and medical markets.

The trend towards increased chip integration and its impact on interfaces to a processor chip will ultimately change the nature of transactions over the interconnect, causing a migration from client/server model to peer-to-peer model. Today's interfaces whose implementation requires a separate chip and are connected to the processor via a PCI bus will be combined on a single chip with multiple interfaces in the future. Eventually, these interfaces will be incorporated with an I/O processor chip. This shift in transaction models will gradually diminish the need for a separate I/O bus like PCI, making the emergence of RapidIO technology especially critical.

Both parallel and serial RapidIO configurations are ideal for embedded and real-time applications where latency, reliability, power consumption, and low cost are high priorities, and eliminates the need to fragment the market along processor bus and board-level interconnect lines.

Parallel RapidIO technology spans the application spectrum from chip-to-chip through board-to-board implementation, specifically optimized for processor bus, chip-to-chip, and mezzanine card interconnects. Its application overlaps the target application areas for HyperTransport and 3GIO, as well as a portion of those for StarFabric and GigaBridge.

Serial RapidIO technology has been optimized for backplane and DSP farm interconnects and can be used to connect multiple computer chassis as a single system. Its application overlaps that of the parallel RapidIO configuration and extends into applications now utilizing FibreChannel, in addition to a portion of those applications targeted by InfiniBand.

RapidIO technology will also provide the backbone for chip-to-chip interconnects on a board. This technology is designed to function as a complement to PCI while there remains a need for local I/O buses.

In all cases, RapidIO is targeted for use within a system where lower latency and simpler implementation are needed to meet latency, power and cost requirements. To connect between systems, a SAN solution such as InfiniBand or FibreChannel should be used in order to achieve greater system isolation at the expense of higher latency.

RapidIO Features and Benefits

Performance

The RapidIO fabric-based architecture connects all the chips and devices within a system directly to each other, allowing for multiple, simultaneous transactions while supporting transmission speeds of 10 Gb/s and beyond. Data paths are 8- or 16-bits wide and data is sampled on both edges of the clock. The LVDS technology utilized by RapidIO has the capability to scale to multi-GHz speeds, and the port width can scale to 16-bits. The availability of large amounts of aggregate bandwidth increases system performance hundreds of times over today's hierarchical bus interconnects.

Low latency (the time lapse between the request for a transaction and the start of that transaction) is achieved through the use of small packet headers and the absence of a software protocol stack. These small headers are organized for fast, efficient assembly and disassembly, while their small size controls overhead.

The RapidIO architecture enables point-to-point determinism with its ability to prioritize the passage order of the most important transactions to maintain Quality of Service (QoS). Multiple message priority levels guarantee that the most critical transmissions reach their destination, even under the most challenging conditions.

Scalability

The RapidIO architecture is a scalable technology that supports a number of system topologies, address maps and transactions to suit a variety of applications. These topologies include point-to-point, ring, star, linked star, mesh and arbitrary fabric topologies, the star configuration being the most common.

In addition, the communications bandwidth of the RapidIO switch fabric can increase proportionally with the number of attached devices. A practically unlimited number of new devices or peripherals can be added to a system designed with RapidIO, with minimal interconnect degradation.

The data rate of each RapidIO link can be implemented in a variety of data widths and clock frequencies. Links of different frequencies can exist within a single system, resulting in the greatest design flexibility and forward compatibility. RapidIO interconnects have the potential to evolve over time to support various segments of the embedded market.

History

RapidIO technology is the result of a joint development effort between Mercury Computers, Inc., and the Motorola Semiconductor Products Sector. Each company originally had its own high-performance interconnect development program; Mercury's was targeted at a new generation of their highly successful RACEway and RACE++ architecture. Each became aware of the other's work through their long-standing supplier-customer relationship and, since objectives were similar, collaboration was logical. Mercury brought both RACEway technology and expertise in multiprocessor computing systems to the joint effort, while Motorola contributed chip architecture expertise and experience in delivering solutions to the networking market. From the beginning, the development of RapidIO technology was envisioned as collaborative development with a broad set of companies. Motorola's customers who manufacture networking equipment were the first to be asked to contribute to the early development of the specification, and their input was incorporated by the time the technology was announced in February 2000. The RapidIO Trade Association was formed shortly thereafter, and the first standard was approved in March 2000 after a full year of intense public review. That specification is publicly available for download from the RapidIO website at www.rapidio.org.

RapidIO

► Technical Brief

Reliability

The RapidIO architecture is the only major fabric interconnect that includes hardware error detection and correction services performed on each individual link in the data transfer path. A separate CRC algorithm is used to detect corruption in the header and data payload, and all control packets are sent redundantly inverted to ensure complete coverage.

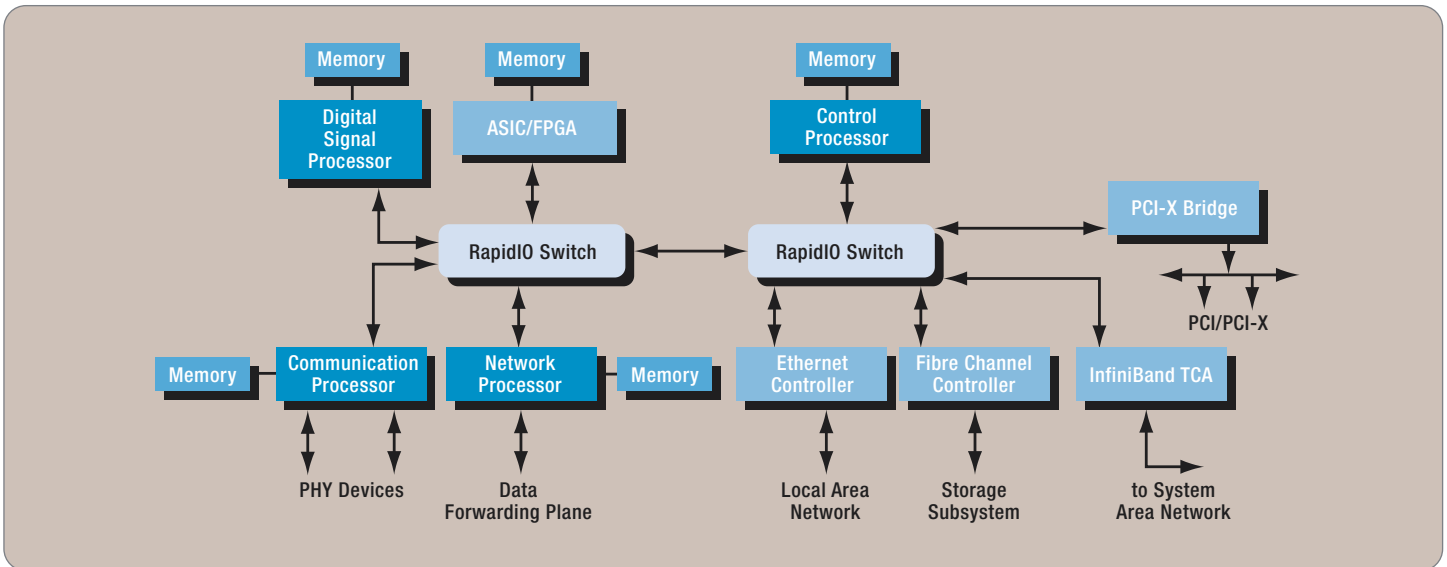
In the event of a non-correctable error, the packet is resent so that no packets are ever lost. Packets can be stomped immediately when a problem is detected, reducing the time required to recover from errors.

The RapidIO architecture also ensures forward progress of information packets. Each packet makes its way through the switch fabric one link at a time. If any particular link is busy, the packet simply waits until the link becomes available rather than retreating.

Compatibility

RapidIO technology offers software compatibility with existing applications on devices with typical load-store architecture, including microprocessor external interfaces and PCI device drivers. Its architecture provides a common connection standard for general-purpose processors, digital signal processors, communication processors, network processors, peripheral devices, and bridges to legacy buses. This architecture supports all needed microprocessor and I/O transactions, making it transparent to existing applications and operating system software and eliminating the need for software vendors to rewrite their core system interface programs. The RapidIO interconnect looks, to software, like a traditional microprocessor and peripheral bus.

For example, this bus technology is easily bridged to PCI, PCI-64 and PCI-X, offering software transparency for PCI-based systems without the need for special device drivers. It is also easily bridged to existing interconnects, such as RACEway, and to SAN interconnects, such as InfiniBand. In addition, RapidIO may be adapted to interface with Ethernet, ATM, USB and the like through the use of protocol chips.



► **Figure 1.** RapidIO architecture is a point-to-point interconnect that bridges to existing systems and networks.

Flexibility

RapidIO is divided into a three-layered hierarchy of logical, transport and physical layers, which provides the flexibility to add to or modify one layer without affecting the others.

The logical layer conducts PCI-like read and write transactions and port-based operations that allow devices to communicate with each other without direct address space visibility. This layer also supports globally shared memory, such as distributed cache coherent memory subsystems and operating system primitives.

The transport (routing) layer distributes packets that contain both source and destination device IDs, allowing up to 64,000 devices in a system. The peer-to-peer architecture of this layer enables distributed control, rather than central control, eliminating the need to pass through a common host. This layer allows variable packet payload sizes up to 256 Bytes, provides support for multiple topologies and offers “true switching” for short, deterministic latency.

The physical layer consists of two sub-sets: parallel and serial. The 8- or 16-bit LP-LVDS parallel interface allows up to 32 Gb/s throughput in each direction. The 1x or 4x LP serial interface – one serial pair or four ganged serial pairs – offers up to 10 Gb/s throughput in each direction, leverages the XAUI electrical specification, includes an 8b/10b-encoded clock, and provides error coverage.

Cost

The RapidIO interconnect is designed with a minimum silicon footprint for low-cost, full-custom ASIC- and FPGA-based designs so that it can be implemented in the corner of a processor, reducing the real estate necessary for interconnect processing. The serial release of this interface enables very low pin counts for extreme power and pin-sensitive designs, such as DSPs.

Together, RapidIO’s compatibility with standard FPGAs, the deployment of a RapidIO interface within a small portion of the modern FPGA device, and the possibility of multiple ports in ASICs or microprocessor implementations, enable fast prototyping, low-cost manufacturing and reduced time-to-market for new products. Ultimately, designers can add multiple RapidIO ports to new I/O chips, providing the performance benefits of a fabric interconnect without incurring the costs associated with adding a dedicated switching chip.

Multiprocessing Support

The RapidIO architecture offers hardware-supported symmetric multiprocessing through an optional distributed shared-memory extension. Distributed shared memory is used pervasively in the computer workstation and server markets, and is becoming more popular in high-performance embedded applications. It is also useful to maintain cache coherency for a single processor in systems with distributed memory controllers.

This architecture supports multiple programming models, such as Numa (with physical addresses), ccNuma and message passing, enabling simultaneous, distributed I/O processing and general-purpose multiprocessing within one single system.

RapidIO


► Technical Brief

Hardware Interoperability Platform (HIP)

Although RapidIO bus technology delivers the increased bandwidth, software independence, fault tolerance and low latency needed by designers to develop next-generation networking equipment in the communications market, this new bus technology introduces new compliance and interoperability challenges that the designer must resolve.


A critical element to the success of an emerging standard is the ability to demonstrate that multi-vendor silicon can interoperate seamlessly, making an interoperability platform essential. RapidIO's hardware interoperability platform (HIP) architecture provides a vehicle to facilitate prototyping by multiple vendors around RapidIO technology. This architecture offers a common environment for silicon vendors to demonstrate interoperability, opening doors to many tools and semiconductor vendors.

This HIP architecture employs a common form factor for switch fabrics and end points, and a common connector and pin assignments for power and RapidIO signal paths, the critical elements that impact interoperability testing. Its architecture consists of a motherboard and a RapidIO plug-in card. The HIP motherboard is intended to provide RapidIO connectivity for RapidIO plug-in cards, as illustrated in Figure 2.



► **Figure 2.** HIP Architecture.

The HIP motherboard complies with the ATX specification (REF3) as it pertains to mounting hole locations, ATX form factor, RapidIO plug-in card slot positions, and power connector type, to ensure that the motherboard can be used within an ATX chassis, if desired, and can utilize ATX power supplies. Location of the power connector is constrained only in that it should be within reach of a standard ATX power supply harness when motherboard and power supply are housed within an ATX chassis. Power is supplied by a standard ATX power supply and RapidIO plug-in cards may utilize the card guides if an ATX chassis is used. Figure 3 shows an example implementation of an HIP motherboard.




► **Figure 3.** Example implementation of an HIP motherboard.

RapidIO plug-in card slots may be used to house PCI or RapidIO plug-in cards. RapidIO plug-in cards occupy the equivalent of two standard PCI card slots due to the real estate required to accommodate the width of the RapidIO connectors used, as illustrated in Figure 3. RapidIO plug-in cards offer easier differential trace signal routing and increased component height clearance to accommodate CPU heat sinks and fans. RapidIO plug-in card slots support two power connectors as well as RapidIO differential signal connectors, and may optionally support an additional PCI connector. Figure 4 illustrates an example implementation of plug-in card slots.

The HIP architecture provides not only a common environment for silicon vendors to demonstrate interoperability, but also a common environment for these vendors to evaluate why silicon does not inter-operate. What happens when a 2 Gigahertz data path between two point-to-point devices fails to communicate? When failures occur, designers must be able to accurately view, characterize, analyze and debug these errors.

This platform, in conjunction with test and measurement equipment, such as logic analyzers, oscilloscopes, probes and test software, allows design engineers to accurately characterize and analyze such failures so they can ensure that their designs comply with communications standards.



▶ **Figure 4.** Example implementation of plug-in card slots.

RapidIO Protocol and Transactions

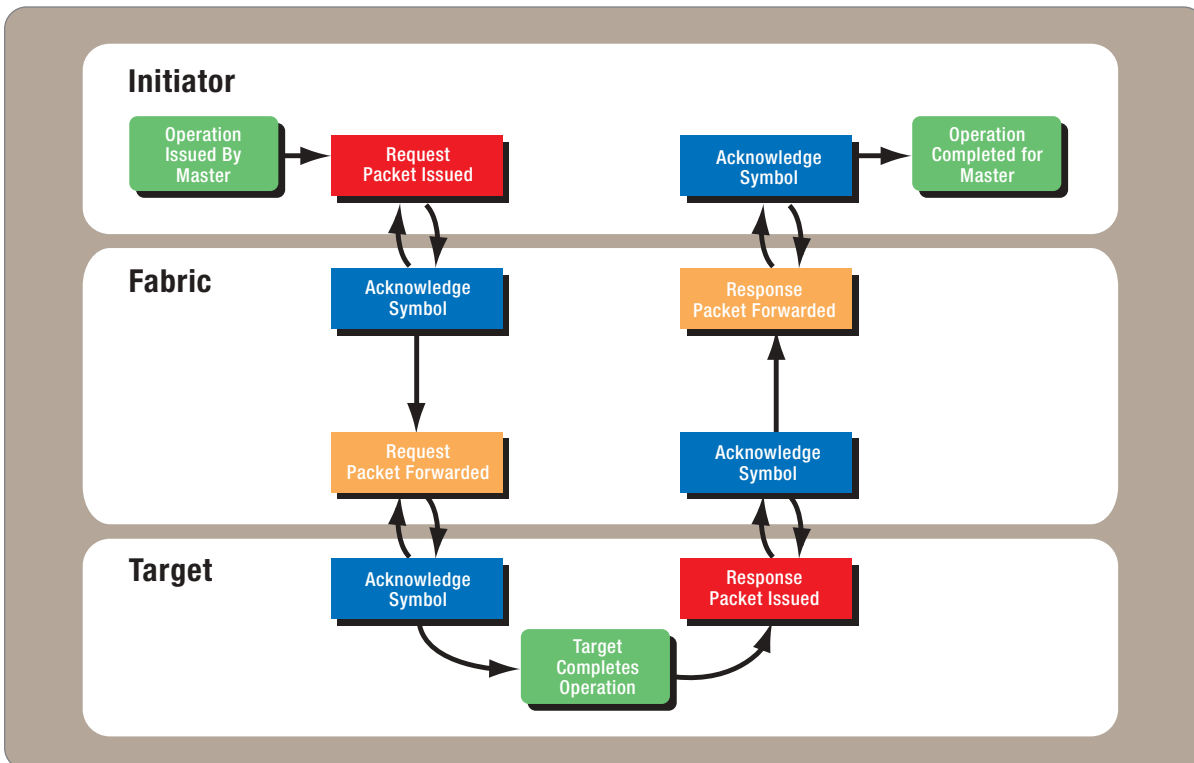
RapidIO end points – initiating and target devices – are rarely connected directly to one another. Most often a message or transaction will pass through an intervening RapidIO fabric. Communication elements – specifically, pairs of request and response packets – will carry the message or transaction between end point devices in the system. Control symbols are used to manage the flow of transactions in the RapidIO physical interconnection layer by performing packet acknowledgement, flow control and maintenance functions.

Flow of a Packet Through a RapidIO System

Figure 5 demonstrates the progress of a packet through a RapidIO system. There may be several fabric devices between the initiating device and the target device, depending on the complexity of the fabric. Packets are held, forwarded and acknowledged between these fabric elements in the same manner. Control symbols may also indicate the detection of an error and a request to retransmit. Data payloads are variable in size, ranging from one to 256 Bytes.

In Figure 5, events occur in the following sequence:

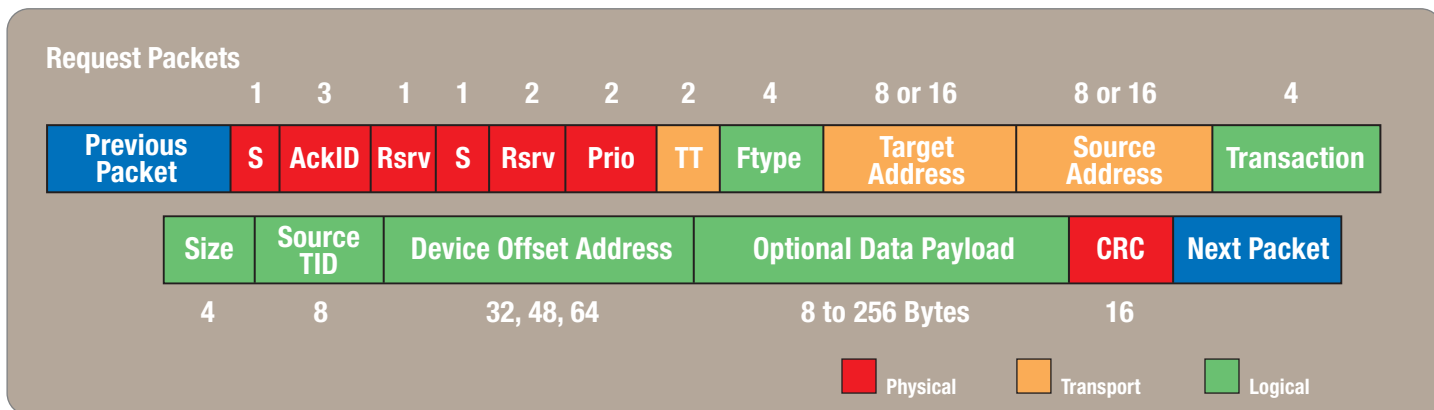
- (1) The initiating device, or master, begins the transaction by issuing a request packet, which is sent to a fabric device.
- (2) The fabric device stores the request packet and sends an acknowledgement control symbol back to the initiator.
- (3) The request packet is forwarded to the target device.
- (4) The target device responds by sending an acknowledgement control symbol back to the fabric device and completing the requested operation.
- (5) Once the requested operation has been completed, the target device sends a response packet back to a fabric device.
- (6) The fabric device stores the response packet and sends an acknowledgement control symbol back to the target device.
- (7) The response packet is forwarded to the initiating, or master, device.
- (8) The master device returns an acknowledgement control symbol to the fabric device.
- (9) The master device now “knows” that the transaction has been completed, and has received response data if present.



► **Figure 5.** Operation sequence.

Packet Construction

Figure 6 illustrates the construction of a request packet. The packet format supports arbitrary widths, especially optimized for 1, 4, 8, 16, and 32-bit wide physical interfaces. The format is partitioned to simplify packet assembly and disassembly in end points. Numbers that appear above the various fields shown are the bit lengths of the relevant fields.



▶ **Figure 6.** Packet formats.

The fields of the packet in Figure 6 are defined as follows:

S:	Defines whether transmission is a data packet or a control symbol.
AckID:	Which packet (of a group) the fabric device or target should acknowledge with a control symbol. RapidIO supports up to 8 unacknowledged packets between two adjacent devices.
PRIO:	Packet priority, used for flow control.
TT:	Type of transport address mechanism used.
Ftype:	Indicates the transaction being requested (see below).
Target Address:	Address to which the packet should be delivered.
Source Address:	Where the packet originated.
Transaction:	Indicates the transaction being requested (see below).
Size:	Encoded transaction size.
srcTID:	Source transaction ID. RapidIO devices may have up to 256 outstanding transactions between two end points.
Device Offset:	Address for memory mapped transactions.
Optional Data Payload:	Intended for “write” operations. These range from one to 256 bytes in size.
CRC:	Intermediate Cyclic Redundancy Check error detection. Only present in long packets.
Optional Data Payload:	Continuation of data payload.
CRC:	Final error detection. Present in both short and long packets.

RapidIO

► Technical Brief

Response packets are similar in construction. The “size” field in the request packet is replaced by a “status” field in the response packet, which indicates whether or not the transaction was successfully completed. The “srcTID” field in the request packet is replaced by a “TargetTID” field in the response packet, which contains the corresponding request packet transaction ID. The “Device Offset Address” field in the request packet is not present in the response packet.

Both request and response packets include “Ftype” and “transaction” fields, which are used to define the type of transaction according to the following table:

Ftype	Class	Transaction Examples	Logical Specification
0, 15	User	User Defined	All
1	Intervention Request	Read from current “owner”	GSM
2	Non-Intervention Request	Read from home, Non-coherent read, I/O read, TLB sync, Atomic	GSM, IOS
5	Write request	Cast-out, Flush, Non-coherent write, Atomic swap	GSM, IOS
6	Streaming Write	Stream Write	IOS
8	Maintenance	Configuration, control, and status register read and write	All
10	Doorbell	In-band Interrupt	MSG
11	Message	Mailbox	MSG
13	Response	Read and write responses	All
3, 4, 7, 9, 12, 14	Reserved		

GSM = Globally Shared Memory Extensions (optional)

IOS = Basic Input/Output System

MSG = Message Passing Extensions

Flow Control

Flow control is defined as a part of the RapidIO physical specification. There are three types of flow control: retry, throttle and credit-based.

The **retry** type of flow control is the most simple mechanism for flow control and is a part of the hardware error recovery. If a receiving device is incapable of accepting a packet, either due to a lack of resources or because an error has been detected, the device may respond with a control symbol that requests retransmission of the packet.

The **throttle** type of flow control allows a device to insert “wait-states” in the middle of packets through use of an idle control symbol. A receiving device can also request that transmission be slowed down, using this idle control symbol.

The **credit-based** type of flow control is used with devices that include transaction buffers or buffer pools, most notably fabric devices. Some control symbols contain a buffer status field; a sender will only transmit a packet when the destination device has sufficient available buffer space in which to store it.

Conclusion

The significant increase in transmission speeds afforded by the RapidIO architecture enables the design of next-generation networking and communications equipment, but with it comes compliance and interoperability issues that design engineers must resolve before implementing RapidIO ports into their designs.

The hardware interoperability platform (HIP) is designed to facilitate prototyping by multiple vendors around the RapidIO technology, enabling them to demonstrate interoperability. Used in conjunction with this platform, test and measurement tools, like logic analyzers, oscilloscopes, probes and test software, gives these designers the tools they need to evaluate the source of failures when they occur.

For information about Tektronix solutions for RapidIO implementation, please visit www.tektronix.com/rapid_io.

RapidIO Implementation

A special backplane will be required in order to implement board-to-board RapidIO communication. No specification has been developed as yet. However, current efforts to standardize generic high-speed serial connections on a backplane apply to serial RapidIO technology, and momentum exists toward leveraging the new serial work in PICMG.



The Integrated Tool Set for RapidIO Implementation

Together, Tektronix world-class logic analyzers, high-performance oscilloscopes and industry-leading probes deliver superior probing, triggering, display and analysis capabilities to enable design engineers to quickly and easily implement RapidIO into their designs.

Contact Tektronix:

- ASEAN / Australasia / Pakistan** (65) 6356 3900
- Austria** +43 2236 8092 262
- Belgium** +32 (2) 715 89 70
- Brazil & South America** 55 (11) 3741-8360
- Canada** 1 (800) 661-5625
- Central Europe & Greece** +43 2236 8092 301
- Denmark** +45 44 850 700
- Finland** +358 (9) 4783 400
- France & North Africa** +33 (0) 1 69 86 80 34
- Germany** +49 (221) 94 77 400
- Hong Kong** (852) 2585-6688
- India** (91) 80-2275577
- Italy** +39 (02) 25086 1
- Japan** 81 (3) 3448-3111
- Mexico, Central America & Caribbean** 52 (55) 56666-333
- The Netherlands** +31 (0) 23 569 5555
- Norway** +47 22 07 07 00
- People's Republic of China** 86 (10) 6235 1230
- Poland** +48 (0) 22 521 53 40
- Republic of Korea** 82 (2) 528-5299
- Russia, CIS & The Baltics** +358 (9) 4783 400
- South Africa** +27 11 254 8360
- Spain** +34 (91) 372 6055
- Sweden** +46 8 477 6503/4
- Taiwan** 886 (2) 2722-9622
- United Kingdom & Eire** +44 (0) 1344 392400
- USA** 1 (800) 426-2200
- USA (Export Sales)** 1 (503) 627-1916

For other areas contact Tektronix, Inc. at: 1 (503) 627-7111

Updated 17 June 2002

For the most up-to-date product information visit our web site at www.tektronix.com



Copyright © 2002, Tektronix, Inc. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.
07/02 HMM/BT 5AW-15975-0