VMICPCI-5579
Economical, High Capacity Reflective Memory (A Real-Time Network Product)

- Real-time network for computers with CompactPCI®, PCI, or PMC interfaces
- Up to 64 Mbyte of Reflective Memory with network data rate up to 13.4 Mbyte/s
- High-speed, easy-to-use fiber-optic network (270 Mbaud serially)
- Data written to memory in one node is also written to memory in all nodes on the network
- Up to 1,000 m between nodes and up to 256 nodes
- Data transferred at 13.4 Mbyte/s without redundant transfer
- Data transferred at 6.7 Mbyte/s with redundant transfer
- Any node on the network can generate an interrupt in any other node on the network or in all network nodes with a single command
- Error detection
- Redundant transmission mode for suppressing errors
- No processor overhead
- No processor involvement in the operation of the network
- D32:D16:D8 memory access
- Single 3U CompactPCI board with optional 6U front panel
- PCI target data bursts supported with 33 Mbyte/s transfer rates
- PCI master DMA controller. DMA transfer rates of 132 Mbyte/s maximum and 22 Mbyte/s sustained.
- Configurable endian conversions for multiple CPU architectures on the network

INTRODUCTION — The VMICPCI-5579 is a member of a high-performance, daisy-chained, fiber-optic network. Data is transferred by writing to on-board global RAM. The data is automatically sent to the location in memory on all Reflective Memory boards on the network.

PRODUCT OVERVIEW — The Reflective Memory concept provides a very fast and efficient way of sharing data across distributed computer systems.

VMIC's VMICPCI-5579 Reflective Memory interface allows data to be shared between up to 256 independent systems (nodes) at rates up to 13.4 Mbyte/s. Each Reflective Memory board may be configured with 32 or 64 Mbyte of on-board RAM. The local RAM provides fast Read access times to stored data. Writes are stored in local RAM and broadcast over a high-speed fiber-optic data path to other Reflective Memory nodes. The transfer of data between nodes is software transparent, so no I/O overhead is required. Transmit and Receive FIFOs buffer data during peak data rates to optimize CPU and bus performance to maintain high data throughput.

The Reflective Memory also allows interrupts to one or more nodes by writing to a byte register. Three separate user-definable interrupts may be used to synchronize a system process, or used to follow any data. The interrupt always follows the data to ensure the reception of the data before the interrupt is acknowledged.

The VMICPCI-5579 requires no initialization unless interrupts are being used or endian byte swapping is desired.

Each node on the system has a unique identification number between 0 and 255. The node number is established during hardware system integration by placement of jumpers on the board. This node number can be read by software by accessing an on-board register. In some applications, this

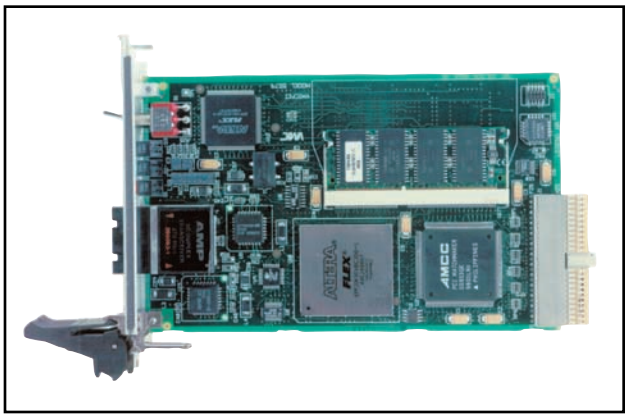

node number would be useful in establishing the function of the node.

LINK ARBITRATION — The VMICPCI-5579 system is a fiber-optic daisy chain ring as shown in Figure 1. Each transfer is passed from node-to-node until it has gone all the way around the ring and reaches the originating node. Each node retransmits all transfers that it receives except those that it originated. Nodes are allowed to insert transfers between transfers passing through.

Ordering Options

Sep. 26, 2002 800-655579-000 B		A	B	C	–	D	E	F
VMICPCI-5579	–			0	–			

A = Memory Options
0 to 3 = Reserved
4 = 32 Mbyte
5 = 64 Mbyte
B = Panel Options
0 = 3U Panel
1 = 6U Panel
C = 0 (Option reserved for future use)

Connector Data

SC Connector

AMP 503948-1 (or equivalent)

Cable Specifications

Fiber-Optic Cable – Multimode; 62.5 Micron core.
Transmitters operate at 1,300 nm at 270 Mbaud.
Maximum attenuation between nodes is 9 dB.
Minimum attenuation between nodes is .5 dB.

Fiber-Optic Cable Assemblies		A	B	C	–	D	E	F
VMICBL-000-F4	–	0			–			

A = Connector Type
0 = Ceramic Ferrule SC Connector
BC = Cable Lengths

02 = 5 ft (1.5 m)	12 = 500 ft (152.4 m)
04 = 25 ft (7.6 m)	16 = 1,000 ft (304.8 m)
05 = 50 ft (15.2 m)	19 = 1,500 ft (457.3 m)
07 = 100 ft (30.4 m)	22 = 2,000 ft (609.7 m)
08 = 150 ft (45.7 m)	25 = 2,460 ft (750.0 m)
09 = 200 ft (60.9 m)	29 = 3,280 ft (1,000 m)
11 = 350 ft (106.7 m)	

For Ordering Information, Call:
1-800-322-3616 or 1-256-880-0444 • FAX (256) 882-0859
E-mail: info@vmic.com Web Address: www.vmic.com

INTERRUPT TRANSFERS — In addition to transferring data between nodes, the VMICPCI-5579 will allow any processor in any node to generate an interrupt on any other node. These interrupts would generally be used to indicate to the receiving node that new data has been sent and is ready for processing. These interrupts are also used to indicate that processing of old data is completed and the receiving node is ready for new data.

Three interrupts are available. The user may define the function for each interrupt. Any processor can generate an interrupt on any other node on the network. In addition, any processor on the network can generate an interrupt on all nodes on the network. Interrupts are generated by simply writing to a VMICPCI-5579 register.

All data and interrupt command transfers contain the node number of the node that originated the transfer. This information is used primarily so the originating node can remove the transfer from the network after the transfer has traversed the ring. The node identification is also used by nodes receiving interrupt commands. When a node receives an interrupt command for itself, it places the identification number of the originating node in a FIFO. Up to 64 interrupts can be stacked in the FIFO. During the interrupt service routine, the identification of the interrupting node can be read from the FIFO. Twenty-one bits of data are also transmitted with the interrupt and stored in the interrupt FIFO.

PCI INITIATOR/DMA CAPABILITIES — The VMICPCI-5579 supports DMA operations. The DMA capability is initiated by a PCI host. After the DMA engine is initialized by a host, the VMICPCI-5579 will request the PCI bus as appropriate and move up to 64 Mbyte as a PCI initiator. This capability removes the CPU from the responsibility of requesting the PCI bus and moving the data itself. This feature is very useful for moving large blocks of data. The PCI architecture ensures that the VMICPCI-5579 does not monopolize the PCI bus during this process. Large DMA blocks will automatically be split into smaller bursts on the PCI bus by the DMA engine. The VMICPCI-5579 can be programmed to issue a PCI interrupt at the conclusion of a DMA transfer or the host can poll the status of the DMA process. The status of the DMA can be accessed from the board, if required. Although the DMA engine can do both DMA reads and DMA writes, they cannot occur simultaneously. The VMICPCI-5579 can burst data onto the PCI bus at a maximum rate of 132 Mbyte/s and sustain 22 Mbyte/s. If the target operates at a slower rate, the PCI handshaking capabilities will throttle the data rate.

ERROR MANAGEMENT — Errors are detected by the VMICPCI-5579 with the use of the error detection facilities of the Hot Link chipset and additional parity encoding and checking. The error rate of the VMICPCI-5579 is a function of the rate of errors produced in the optical portion of the system. This optical error rate depends on the length and type of fiber-optic cable. Assuming an optical error rate of 10^{-12}, the error rate of the VMICPCI-5579 is 1.6×10^{-10} transfers/transfer.

However, the rate of undetectable errors is less than 2.56×10^{-20} transfers/transfer. When a node detects an error, the erroneous transfer is removed from the system and a PCI bus interrupt is generated, if enabled.

The VMICPCI-5579 can be operated in a redundant transfer mode in which each transfer is transmitted twice. In this mode of operation, the first of the two transfers is used unless an error is detected in which case the second transfer is used. In the event that an error is detected in both transfers, the node removes the transfer from the system. The probability of both transfers containing an error is 2.56×10^{-20}, or about one error every 1,470,000 years at maximum data rate.

ENDIAN CONVERSIONS — Data lane steering can be configured in a Control Register to allow CPUs of different architectures to communicate. Byte swap, Word swap, and Byte-Word swap options are available.

PROTECTION AGAINST LOST DATA — Data received by the node from the fiber-optic cable is error checked and placed in a receive FIFO. Arbitration with accesses from the PCI bus then takes place, and the data is written to the node's RAM and to the node's transmit FIFO. Data written to the board from the PCI bus is placed directly into RAM and into the transmit FIFO. Data in the transmit FIFO is transmitted by the node over the fiber-optic cable to the next node.

The product is designed to prevent either FIFO from becoming full and overflowing. It is important to note the only way that data can start to accumulate in FIFOs is for data to enter the node at a rate greater than 13.4 Mbyte/s or 6.7 Mbyte/s in redundant mode. Since data can enter from the fiber and from the PCI bus, it is possible to exceed these rates. If the transmit FIFO becomes half-full, a bit in the Status Register is set. This is an indication to the node's software that subsequent **writes** to the Reflective Memory should be suspended until the FIFO is less than half-full. If the half-full indication is ignored and the transmit FIFO becomes full, then writes to the Reflective Memory will be acknowledged with a STOP*. No data will be lost.

NETWORK MONITOR — There is a bit in a Status Register that can be used to verify that data is traversing the ring (that is, the ring is not broken). This can also be used to measure network latency.

VMIC offers single fiber cable assemblies and adapters that are compatible with the VMICPCI-5579 in length ranging from 1.5 to 1,000 m. These cable assemblies are U.L./NEC-rated OFNP and have a 2.5 mm SC-style connector on each end.

SPECIFICATIONS

Memory Size: 32 or 64 Mbyte

PCI Transfer Rate: 33 Mbyte/s as a Target, 22 Mbyte/s as DMA Initiator

TRANSFER SPECIFICATION

Transfer Rate: 13.4 Mbyte/s (longword accesses) without redundant transfer
6.7 Mbyte/s (longword accesses) with redundant transfer

MEMORY CONFIGURATION

The host system automatically allocates to the VMICPCI-5579 a memory space on the CompactPCI bus which is equal to the board memory option. All memory contained on the VMICPCI-5579 is fully read/write accessible by the PCI host system with the exception of the first 40 (HEX) bytes which have been replaced with status and control registers.

SOFTWARE DRIVER

VxWorks and Windows NT® drivers are available. The VMIPCI/SW-RFM1 Network and Shared Memory Driver provides an applications program with three convenient methods for exchanging data among hosts connected to the same Reflective Memory network:

1. Programmed I/O (Peek and Poke) - An applications program can treat the memory on the Reflective Memory device as ordinary memory in which the program can use ordinary load and store accesses.
2. Direct Memory Access (DMA) - On systems where the performance penalty for individual bus accesses is unacceptably high, the DMA feature is available to transfer data in variable-sized blocks.
3. TCP/IP Protocols - Full support for the internetworking TCP/IP protocols allows peer-to-peer communication between applications without the need to design and implement custom communications protocols.

INTERCONNECTION

Cable Requirements: Two fiber-optic cables

Cable Length: 1,000 m maximum between nodes

Configuration: Daisy-chain ring up to 256 nodes

Power Requirements: 3 A maximum at +5 VDC
2 A maximum at 3.3 VDC

PHYSICAL/ENVIRONMENTAL

Temperature Range: 0 to 65 ˚C, operating with forced air cooling. -40 to 85 ˚C, storage.

Relative Humidity: 20 to 80 percent, noncondensing

DATA TRANSFERS

Data written into the Reflective Memory is broadcast to all nodes on the network without further involvement of the sending or receiving nodes. Data is transferred from memory locations on the sending nodes to corresponding memory locations on the receiving nodes.

A functional block diagram of the VMICPCI-5579 is shown in Figure 2.

TRADEMARKS

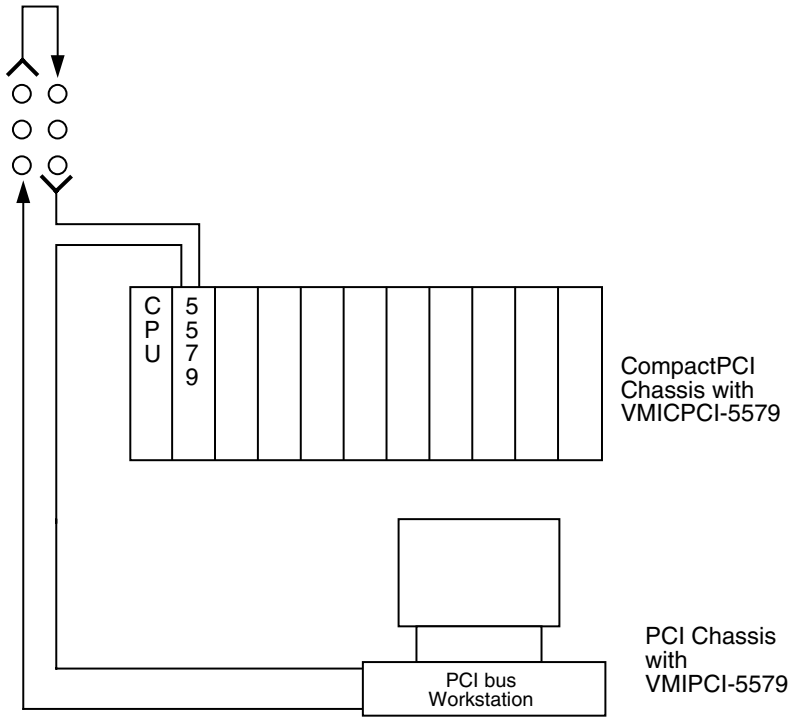

Figure 1. Network Example Using Reflective Memory System

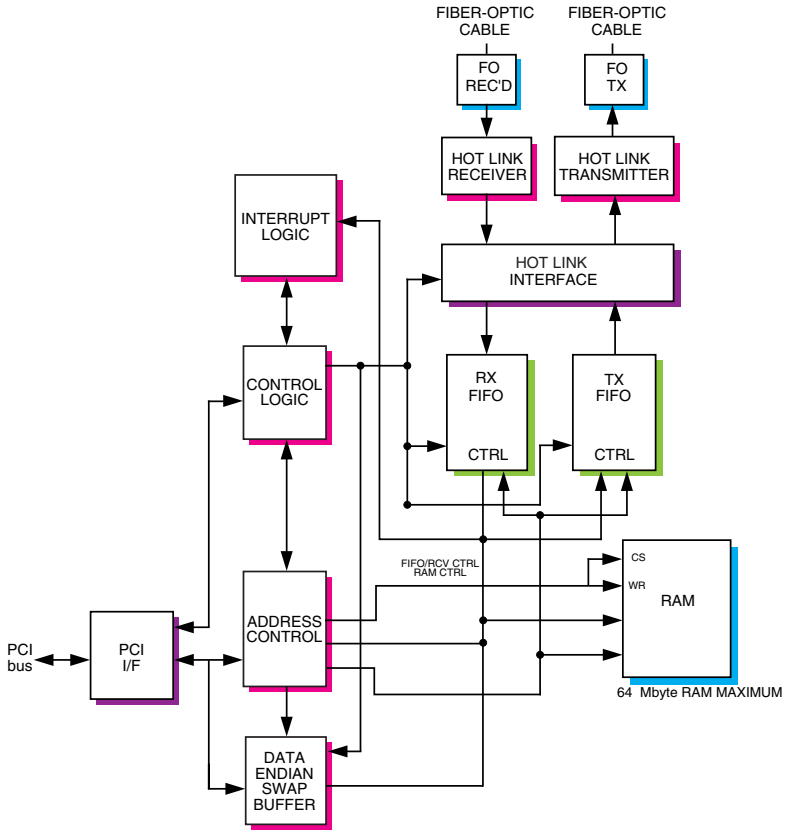

Figure 2. VMICPCI-5579 Functional Block Diagram