



The VMIPCI-5587 issues network interrupts to one or more nodes by writing to a byte/word/longword register. Three separate, user-definable interrupts may be used to synchronize a system process, or to signal the transfer of data. The data and interrupts are always transferred over the network in the same order as issued so the reception of data is ensured before the interrupt is acknowledged.

Basic operation of the VMIPCI-5587 requires no initialization unless endian byte swapping is desired. The interrupts and the extended features, on the other hand, will require software interaction.

Each node on the network has a unique identification number between 0 and 255. The node number or ID is set prior to board installation by installing and/or removing a group of jumpers. This node number can be read by software by accessing an on-board register.

In order to achieve an aggregate throughput rate of 29.5 Mbyte/s, the node processors must be capable of writing to their VMIPCI-5587 board memory at that rate.

**LINK ARBITRATION** — The network including a VMIPCI-5587 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 original node. Each node retransmits all transfers that it receives except those that it originates.

**INTERRUPT TRANSFERS** — The VMIPCI-5587 provides three network interrupts. Any processor can generate an interrupt on any other node on the network. In addition, any processor can generate an interrupt on all nodes on the network with a single register write.

In response to this interrupt register write, the sending VMIPCI-5587 issues a special packet over the network, which contains the command strobe, the sender node ID, the destination node ID, three flags indicating the interrupt type, and 5 to 21 bits of data. When a receiving node detects the proper combination of destination node ID and command strobe, it stores the sender node ID, the flags, and the data in one of three 63 location-deep FIFOs. The three FIFOs correspond to the three interrupts. Upon storing this information in a FIFO, the receiving node issues an interrupt to the local processor if it has been software-enabled. The 5 to 21 bits of data stored in the FIFO is user-definable and typically is treated as an interrupt vector. As part of an interrupt service routine, the local processor reads this information out of the FIFO and acts accordingly. (Note: The VMIPCI-5587 interrupt features are backward-compatible with the VMIxxx-5588 series. However, the VMIxxx-5588 series stores only the sender node ID and not data or flags in its interrupt FIFOs.)

**PCI INITIATOR/DMA CAPABILITIES** — The VMIPCI-5587 supports DMA operations. The DMA sequence is initialized by a few control register writes to the VMIPCI-5587 by the host. Therefore, the VMIPCI-5587 becomes a PCI initiator and moves the specified block of data up to 64 Mbyte without further CPU attention. The PCI architecture ensures that the VMIPCI-5587 does not monopolize the PCI bus and causes the VMIPCI-5587's

DMA engine to automatically split large blocks in small bursts. The VMIPCI-5587 can be programmed to issue a PCI interrupt upon completion of DMA process. Although the DMA engine can do both DMA reads and DMA writes, they cannot occur simultaneously. The VMIPCI-5587 can burst data on to the PCI bus at a maximum rate of 44 Mbyte/s (three PCI clocks).

**ERROR MANAGEMENT** — Errors are detected by the VMIPCI-5587 with the use of the error detection facilities of the Fibre Channel encoder/decoder and additional interlaced parity encoding and checking. The error rate of the VMIPCI-5587 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 VMIPCI-5587 is  $1.3 \times 10^{-10}$  transfers/transfer.

However, the rate of undetectable errors is less than  $1.64 \times 10^{-20}$  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 VMIPCI-5587 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  $1.64 \times 10^{-20}$ , or about one error every 317,855 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 SRAM and to the node's transmit FIFO. Data written to the board from the PCI bus is placed directly into SRAM and into the transmit FIFO. Data in the transmit FIFO is transmitted by the node over the fiber-optic cable to the next node. Data could be lost if either FIFO were allowed to become full.

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 29.5 Mbyte/s or 14.8 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.

If the receive FIFO is allowed to become over half-full, there is a danger the receiver FIFO may overflow resulting in

data loss. In order to prevent this situation, all PCI writes will be acknowledged by a STOP\* until the receive FIFO is less than half-full.

**MEMORY BLOCK PROTECT** — The VMIPCI-5587 has the ability to selectively prevent other nodes from writing to a segment of on-board memory. The block protection feature segments the 64 Mbyte of on-board memory into 256 blocks of 256 Kbyte each. The block protection circuit can prevent any one or more of the 256 node IDs from writing to any one or all 256 blocks of on-board memory. To operate, the host must initialize an SRAM array that represents the 256 segments and 256 node IDs and then must enable the block protection feature. If the block protection feature is activated and an access fault occurs, the VMIPCI-5587 can be programmed to issue a PCI interrupt to the host processor. The host can then read a status register which contains the node ID and segment number of the last invalid write.

Although the block protection prevents writes to on-board RAM where required, it does not alter the data packet as it passes through to the next node in the network. In this manner, several groups of nodes can share the same daisy-chained link and yet operate relatively independently where desired.

**DIAGNOSTIC FEATURES** — The VMIPCI-5587 contains several control and status registers, which may serve to indicate the configuration and status of various sections of the board. Many of the registers have addresses and bit assignments that are identical to those of the VMIPCI-5588. One such status register contains an own data flag which can be used to verify that data is traversing the ring (that is, the daisy-chain ring is not broken). The own data flag can also be used to measure network latency.

The VMIPCI-5587 contains several features, which are not present in prior VMIC products such as the VMIXxx-5588 series. They are intended to aid in both local and network-level diagnostics. Those features are described below:

The Monitor Mode is one such feature. Once the monitor mode is activated by a control register write, the host will be inhibited from writing to on-board memory and from writing data to other nodes. Access of the on-board memory by other nodes is unaffected. The host will only be able to monitor network activity by reading the on-board memory.

The Disconnect Mode is another diagnostic feature. Once the Disconnect Mode is activated through a control register write, the on-board memory will be disconnected from the network. In this mode, the host can perform on-board memory tests without affecting other nodes on the network. At the same time, other nodes can pass data through the board without affecting local memory.

The Upstream Node ID register indicates the node ID of the immediate upstream neighbor in the daisy-chain. The upstream node is the one connected to a particular board's fiber-optic receiver. With this information and software routine that utilizes the interrupts with data, the entire new network can mapped.

The Initiator Node List is contained within a small SRAM in the form of a bit pattern which details all nodes (by ID) that have initiated data packet write on the network. The Initiator Node List SRAM can be overwritten and reinitialized.

**MEMORY CONFIGURATIONS** — The VMIPCI-5587 provides memory size options of 1, 4, 8, 16, 32, and 64 Mbyte. The 1, 4, and 8 Mbyte options are offered to ensure compatibility in existing networks of VMIXxx-5588 series products that operate at those memory sizes. As with any typical PCI bus product, the host system automatically allocates a memory space on the PCI bus, which is equal to the board memory option. All the on-board memory contained on the VMIPCI-5587 is fully read/write accessible by the host system with the exception of the first 40 (HEX) bytes. The first 40 (HEX) memory bytes have been replaced with control and status registers and several of the bits within those registers are either read-only or write-only.

To correctly pass data beyond the 16 Mbyte address range through the network, all nodes must be of the VMIXxx-5587 type. Although the VMIPCI-5587 has maintained backward-compatibility with VMIXxx-5588 series products, those VMIXxx-5588 products do not detect the address bits beyond the 16 Mbyte range, nor will they pass these address bits on down the network. The VMIPCI-5587 does contain a jumper, which, if installed, ensures that only the first 16 Mbyte of memory is passed through the network. With the jumper installed, memory beyond 16 Mbyte space remains accessible by the host.

**CONTROL AND STATUS REGISTERS** — The first 40 (HEX) byte memory locations have been substituted by control and status registers (CSRs). For the standard features common to both the VMIPCI-5587 and the VMIPCI-5588, the CSR addresses and bit assignments are identical. In most instances, a VMIPCI-5588 can be substituted with a VMIPCI-5587 with little or no software impact. The extended features of the VMIPCI-5587 have been assigned to previously undefined register addresses.

## SPECIFICATIONS

**Memory Size:** 1, 4, 8, 16, 32, or 64 Mbyte

**Transmit FIFO Size:** 1 or 4 K packets deep

**PCI Transfer Rate:** 44 Mbyte/s as a target or as DMA initiator

**Network Transfer Rate:** 29.5 Mbyte/s (longword accesses) in nonredundant transfer mode  
14.8 Mbyte/s (longword accesses) in redundant transfer mode

**Cable Requirements:** Two fiber-optic cables (see cable specifications on page 1 of this document)

**Cable Length:** 1,000 ft between nodes with the ST-style, VMIXxx-5588-compatible transceiver (62.5/125 µm MMF-type cable)

820 ft between nodes with the SC-style, industry-standard Gigabit transceiver (62.5/125 μm MMF-type cable)

**Network Configuration:** Daisy-chained ring of up to 256 nodes

**Power Requirements:** 5.0 A maximum at 5.0 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

**SOFTWARE DRIVERS**

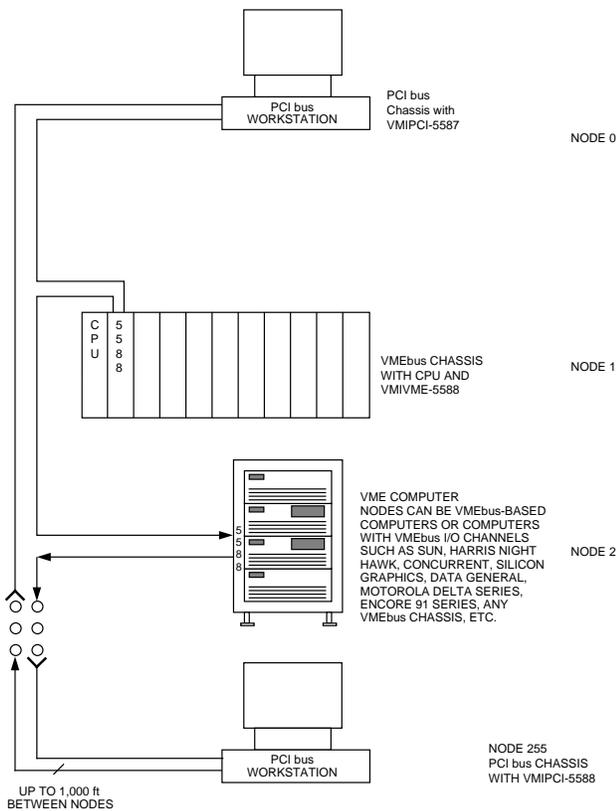
VxWorks and Windows NT drivers are available. The VMIPCI/SW-RFM1 Network and shared memory driver provides an applications program with three

convenient methods for exchanging data among hosts connected to the same Reflective Memory network:

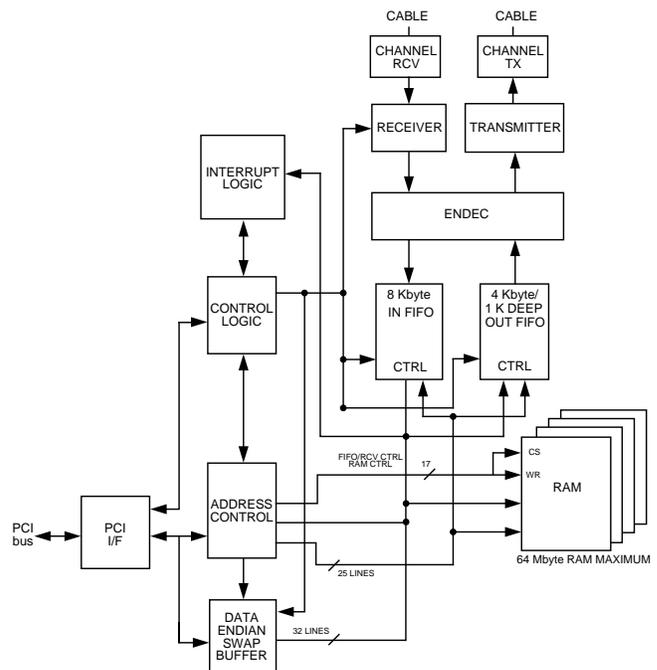
1. Programmed I/O (Peek and Poke) – An applications program can treat the memory on the Reflective Memory device as ordinary memory in which the program can use ordinary load and store accesses.
2. Direct Memory Access (DMA) – On systems where the performance penalty for individual bus accesses is unacceptably 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.

**TRADEMARKS**

RTnet is a trademark and the VMIC logo is a registered trademark of VMIC. Windows NT is a registered trademark of Microsoft Corporation. Other registered trademarks are the property of their respective owners.



**Figure 1. Network Example Using Reflective Memory System**



**Figure 2. VMIPCI-5587 Functional Block Diagram**