500+ Mbps Data Acquisition & Recording using COTS Equipment

Richard O'Connell
Myriad Logic Inc.
1109 Spring Street
Silver Spring MD, 20910
Phone +1-301-588-1900; Fax: +1-301-588-0605
Email: roc@myriadlogic.com

Presented at:
THIC General Meeting
April 23, 1997
Tysons Corner VA.

Introduction

For many years, system designers have employed standards-based, commercial-off-the-shelf (COTS) equipment to lower development cost, speed development cycles, take advantage of new technology, and reduce risk. Usually, this COTS equipment is based on an established system bus standard such as VME or PCI. While VME and PCI buses do provide considerable bandwidth they are not particularly well suited to dealing with the deterministic performance required by high-rate (greater than 120 Mbps) real-time data streams.

Typically, deterministic performance is provided by dedicating secondary data paths (in addition to the system bus) for high-rate, board-to-board data transfer. In the past these secondary data paths were implemented in various ways such as VME Subsystem Buses (VSBs), Digital Signal Processor (DSP) links, and Front Panel Data Ports (FPDP); however, these mechanisms often lacked the bandwidth or flexibility to handle system throughput rates in excess of 200 Mbps. Recently the ANSI standard RACEway has emerged as a high-performance, standards-based secondary data path for VME systems. Using RACEway it is now possible to construct systems with system level throughputs up to 1 Gbit/second using standards-based COTS equipment.

RACEway expands the internal bandwidth of standard VME systems by allowing a board to establish a 160 Mbyte/second point-to-point connection to another board. The RACEway Interlink dynamically manages these connections, allowing many independent 160 Mbyte/second channels to operate simultaneously in the same VME chassis. In this architecture, internal bandwidth scales with the number of boards in the system, a significant advantage over a multi-board shared bus architecture. Additionally, the system designer can guarantee that bandwidth will be available between two critical real-time functions, which is not always the case in systems that rely on shared buses for data transfer. These features make RACEway particularly well adapted to real-time, high-rate applications where a deterministic data flow can be pipelined through one or more functional elements.

This paper discusses how COTS RACEway equipment can be used to acquire, record, and communicate real-time data streams ranging from 200 to 500+ Mbps. First, an overview of RACEway technology is provided including a review of RACEway COTS products available from more than a dozen vendors. Then, specific examples of high-throughput real-time systems that incorporate RACEway and high-performance instrumentation recorders (ID-1 and DCRSi) are discussed. Finally, two specific RACEway benchmarks are provided, indicating system-level throughput in the 1 Gbit/second range is achievable.

RACEway - A Switched Fabric Interconnect for VME Environments

The RACEway specification (ANSI/VITA 5-1994, see references) defines a standard usage of the VME P2 connector's A and C rows ("User Defined" pins in the VME specification) to provide an additional high-speed mechanism for transferring large volumes of data between boards. RACEway uses these pins to implement a dynamically switched, point-to-point interconnect system, also known as a switched fabric. Since RACEway uses only P2 rows A and C, it can operate concurrently with VME or VME64 operations. In fact, since all RACEway connections are point-to-point, RACEway operates more reliably with VME64 than other standard P2 interconnects that are bused and can create cross-talk problems (such as VSB).

The central element of a RACEway interconnect system is the Interlink. The Interlink is a crossbar switch that establishes point-to-point connections between boards in a VME/RACEway environment. Physically, the Interlink is an active P2 overlay that mates to the rear of a VME backplane (similar to VSB overlays). The slots encompassed by the Interlink determine the boards that may participate in RACEway connections. Interlinks can encompass a variable number of slots; for example, Cypress Semiconductor offers 4, 8, and 16 slot RACEway Interlink modules enabling a variety of system-level configurations.

RACEway Interlink modules are generally constructed from one or more six-port crossbar switches. Each switch is a single chip that fully implements the RACEway Interlink standard. Interlink overlays spanning up to six slots can be constructed with a single chip. Interlink overlays spanning more than six slots are constructed by cascading six-port crossbars. While cascading crossbars is an effective mechanism for expanding Interlink modules, it does introduce additional latencies that must be considered in the overall system design. Figure 1 illustrates the RACEway switched fabric in a VME system.

Figure 1 - The RACEway Switched Fabric in a VME Environment

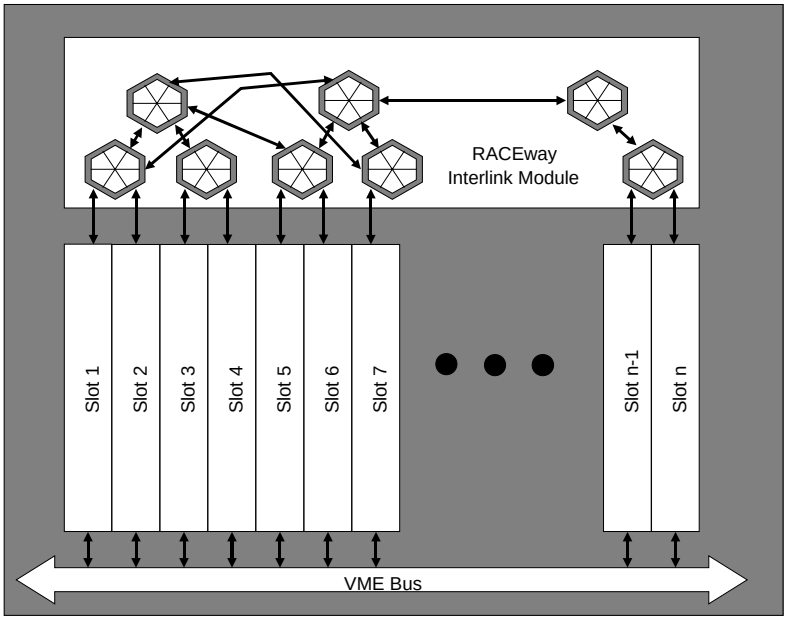

A RACEway master board wishing to transfer data to/from a RACEway slave board issues a request to the fabric by writing a routing word and slave address. The Interlink module receives this information and routes it to the requested slave board. The slave board determines whether a read or write operation has been requested and signals when it is ready to accomplish the data transfer. The Interlink manages actual data transfer on a 2 kbyte burst basis. When both the sender and receiver are able to accomplish a 2 kbyte burst, the Interlink allows the burst to occur. RACEway bursts always occur at 160 Mbytes/second (32 bit synchronous transfers at 40 MHz). The Interlink is capable of interleaving bursts from several RACEway masters to a single RACEway slave. When two or more RACEway masters attempt to simultaneously access multiple-megabyte data blocks to a single RACEway slave, the Interlink arbitrates accesses on a 2 kbyte burst basis. The Interlink also recognizes priorities asserted by various RACEway masters and arbitrates requests accordingly.

The RACEway specification also allows for adaptive routing, split reads, broadcasts, and multicasts. Although these features may be interesting, they are not central to the use of RACEway as a real-time architecture. Additionally, all RACEway-Ready equipment is not required to support these features, so it may not be prudent to design a system assuming these features will be available.

RACEway Applied to Real-Time Systems

The real-time system designer can take advantage of RACEway's multiple data paths by pipelining the data flow through the system. Figure 2 illustrates a typical RACEway-based data acquisition system. In this example, a rotating set of three RACEway memory boards is used to

buffer high-rate, real-time data. As data are acquired, they are written to one memory board. When that memory is filled, the acquisition switches to the second memory, while a data capture board extracts data from the first memory for recording on high-rate tape or Fibre Channel RAIDs. When the second memory is full, acquisition switches to the third memory, data capture switches to the second memory, and signal processors begin operating on data in the first memory. The signal processors can also use RACEway to pass data to additional signal processing boards or a high-performance workstation/supercomputer via a HIPPI or Fibre Channel RACEway board. In this example, the VME bus is used for control only.

The above example illustrates a nine-board system with an internal bandwidth of 640 Mbytes/second (via RACEway) and 80 Mbytes/second (via VME64). This internal bandwidth is considerably higher than can be achieved with industry-standard bus systems (133 and 266 Mbytes/second for 32 and 64 bit PCI, for example). While these numbers are theoretical maximums, systems can be constructed today using COTS equipment that sustain multiple, 100+ Mbyte/second transfers between boards.

While very high internal bandwidth figures are impressive, it is the scalability and determinism of RACEway channels that are the major advantage in real-time systems. In systems where data are passed from one element to the next (such as the example above), scalability is vital in achieving high throughput. In a standard bus architecture, the example system would require all

Figure 2 - Real-Time RACEway System Provides Multiple 160 Mbyte/second Channels

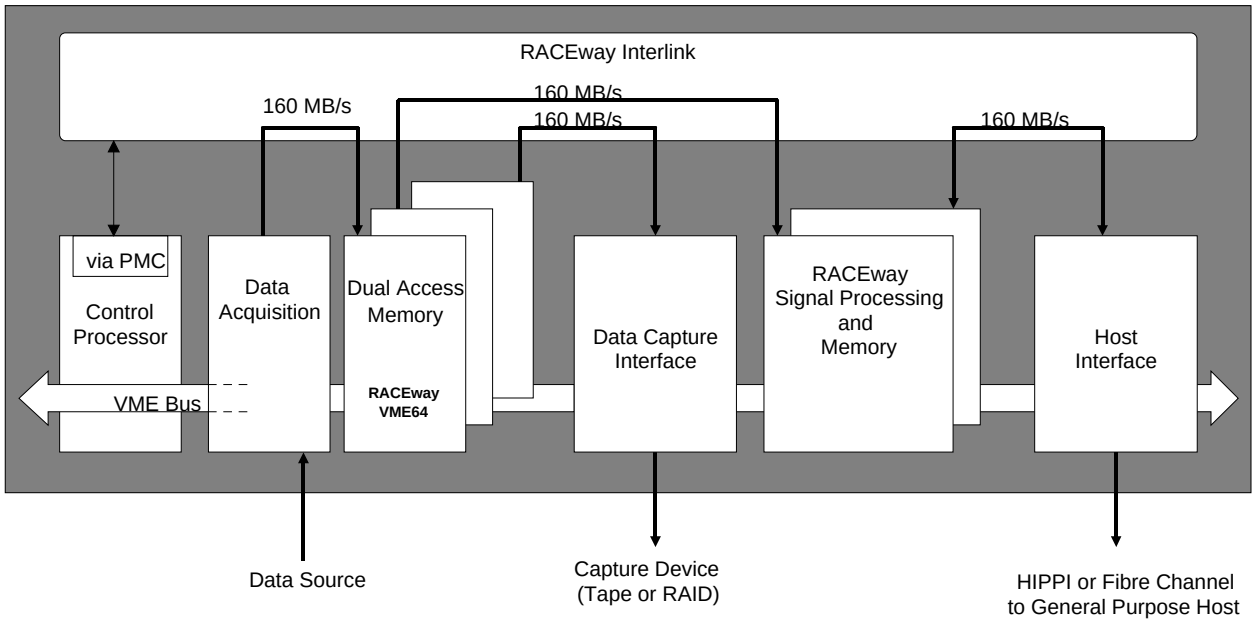

boards to compete for bandwidth on one shared bus. Since the data pass from one element to another four times, the total system throughput in a bused implementation could be no more than one quarter of the maximum achievable transfer rate. In a VME64-only implementation, it

would be very difficult to sustain even a 10 Mbyte/second input rate, while the RACEway extension to VME can realistically sustain 100 Mbytes/second.

Deterministic performance is the second major feature of RACEway important to real-time systems. Individual RACEway links can be dedicated to real-time data paths, insuring that the required bandwidth will be available when needed. In a bused system, the single data pathway may be occupied with a sequence of high-speed bursts or by relatively slow control messages. In typical VME64 systems, latencies due to bus contention can be one or more milliseconds. These latencies do not occur in a RACEway architecture.

RACEway Building Blocks

Widespread industry support has been one of the keys to the success of VME. A large community of vendors offers a plethora of products, many of which are tailored to real-time systems. Since RACEway is fully compatible with VME, it naturally benefits from the great variety of VME products. More importantly, since the adoption of RACEway as an ANSI standard in June 1994, a growing number of manufacturers are supporting RACEway-ready versions of their VME products. COTS RACEway board-level products are available for data acquisition, digital signal processing, recorder interfaces, high-speed communication channels, and a variety of other functions. There are also chip-level products and foundation boards available to assist in the development of application-specific RACEway boards. Table 1 provides a partial list of advertised RACEway products.

Compatibility with VME is a very important feature of RACEway. It enables use of standard VME equipment such as system enclosures (both standard and ruggedized), high-performance single-board computers with standard operating systems (typically used as control elements), and lower performance equipment (such as IRIG time-code generators/readers). Additionally, VME compatibility protects the developer's corporate investment in VME technology, allowing re-use of application-specific VME modules and capitalizing on staff's knowledge and familiarity with VME.

The variety of RACEway-ready products currently available coupled with the ability to integrate these products with standard VME equipment enables the system designer to configure a complex, high-performance real-time system using off-the-shelf products. In fact, the range of VME/RACEway off-the-shelf products is broader than for any competing technology.

Table 1 - RACEway Products

Type	Description	Vendor*
Analog & Digital Input/Output	Analog to 150 MSamples/second	Apcom/Celerity Systems
	Analog to 400 MHz, Digital Receivers	Echotek
	Analog, Digital, Serial to 30 MSamples/second	Pentland
	Common Data Link	Myriad Logic (July 1997)
Memory Boards	128 Mbyte to 2 Gbyte per board	Micro Memories
Digital Signal Processors	Intel I860	Mercury Computer Systems
	Power PC	Mercury Computer Systems
	SHARC	Mercury Computer Systems
	FPGA Preprocessors & FFT	Catalina Research
	TI 320C40 & C80	Mizar
Single Board Computers (SBC)	Various via PMC sites on SBC	Cypress (PMC-RACEway bridge)
Communication Channels	HIPPI (100 Mbyte/sec)	Myriad Logic
Fiber Optic Communications	Serial HIPPI (100 Mbyte/sec)	Myriad Logic
	500 Mbaud	Atlantic Aerospace
RAID Interfaces	SCSI	VMetro foundation with PMC
	SCSI2/Ultra SCSI (40 Mbyte/sec)	Mercury
	Fibre Channel (50 Mbyte/sec)	Atlantic Aerospace
	Fibre Channel (100 Mbyte/sec)	Myriad Logic (June 1997)
		VMetro foundation with PMC
Instrumentation Recorder Interfaces	Ampex DCRsi (30 Mbyte/sec)	Myriad Logic
	Schlumberger ID-1 (60 Mbyte/sec)	Myriad Logic
	Sony ID-1 (64 Mbyte/sec)	Myriad Logic
Foundation Boards	PMC Carrier	VMetro
	General I/O Motherboard	Myriad Logic
Interlink Modules	4, 8, 16 slot	Mercury Computer Systems
		Cypress
Chip-sets	FIFO to RACEway chipsets	Mercury Computer Systems
		Cypress
Bridges	PMC	Mercury Computer Systems
	RACEway/PCI	Cypress

* Contact information for vendors is provided at the end of the article.

The HRS-2000 - Workstation/Supercomputer Controller for 500+ Mbps ID-1 Recorders

The HRS-2000 is an RACEway-based controller used to connect Sony DIR-1000 or Schlumberger DV6000 ID-1 recorders to standard workstations and supercomputers such as the Silicon Graphics Challenge and Origin computers. The simple, three VME/RACEway board construction (Figures 3 and 4) allows the host interface, memory configuration, and recorder options to be changed by simply installing different VME/RACEway modules. Currently, ANSI standard HIPPI and Serial HIPPI connections to the host computer each capable of operating at rates up to 100 Mbytes/second (800 Mbps) are supported. In the third quarter of 1997, 100 Mbyte/second Fibre Channel connections will also be available.

The HRS-2000 uses COTS RACEway memory boards (MMI 6490D) to implement a variable rate buffer. Therefore, the HRS-2000 can be used with non-buffered ID-1 recorders. Typically, the HRS-2000 is configured with 256 Mbytes of memory for applications using the DIR-1000 or DV6400 recorders and with 512 Mbytes for the 512 Mbps DIR-1000H or 520 Mbps DV6800. Memory can be expanded to several Gbytes to support specific host requirements.

RACEway technology is key to the operation of the HRS-2000. Using a single memory module, a 64 Mbyte/second input stream and 64 Mbyte/second output stream are supported simultaneously, requiring a 128 Mbyte/second sustainable internal bandwidth. This internal bandwidth cannot be achieved with current VME64 or 32 bit PCI systems. Additionally, with multiple memory boards, the HRS-2000 can easily support a 100 Mbyte/second input stream simultaneously with the 64 Mbyte/second output stream, demonstrating the scalability of RACEway.

Figure 3 - HRS-2000 - Standard Computer Controller for the 512 Mbps Sony DIR-1000H

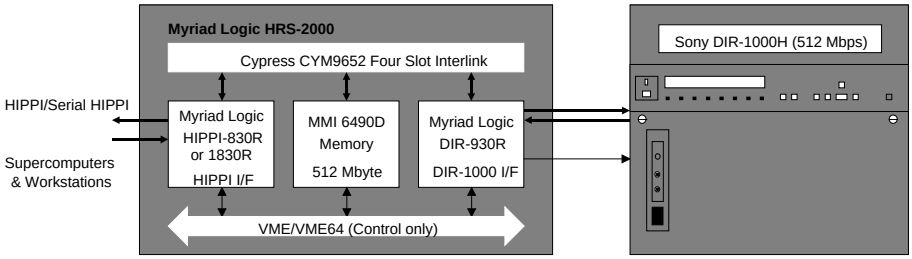

Figure 4 - HRS-2000 Configured for the 520 Mbps Schlumberger DV6800

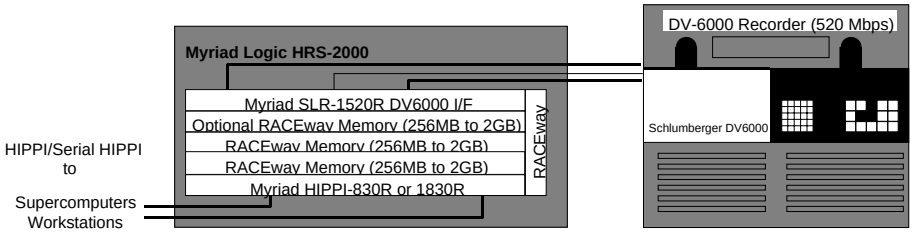

Capture, Processing, and Recording of Aggregate 500 Mbps to 2 Gbps Data Streams

The scalability of RACEway makes it an ideal architecture for handling multiple high-rate data streams. Two examples of such systems are a high-speed data acquisition system under construction at government laboratories, and a high-rate analog data acquisition system.

High-speed Data Acquisition System

The High-speed data acquisition system (Figure 5) is an airborne configuration that acquires data from 6 concurrent 30 Mbyte/second data streams, storing the data in COTS VME/RACEway memory boards. Selected portions of the data are read from the memories and recorded on an Ampex DCRsi. The system uses six 512 Mbyte memory boards (using six slots), for a total of 3 Gbytes of solid-state buffering. By using available 2 Gbyte memory boards, the total buffering can be easily expanded to 12 Gbytes. Furthermore, up to 8 more memory boards can be included in the system, taking the total buffering capacity to 28 Gbytes. The system makes use of existing ruggedized enclosures to provide the shock, vibration, and EMI/EMC protection required for airborne applications.

Figure 5 - RACEway System that Acquires Data at 180 Mbytes/second

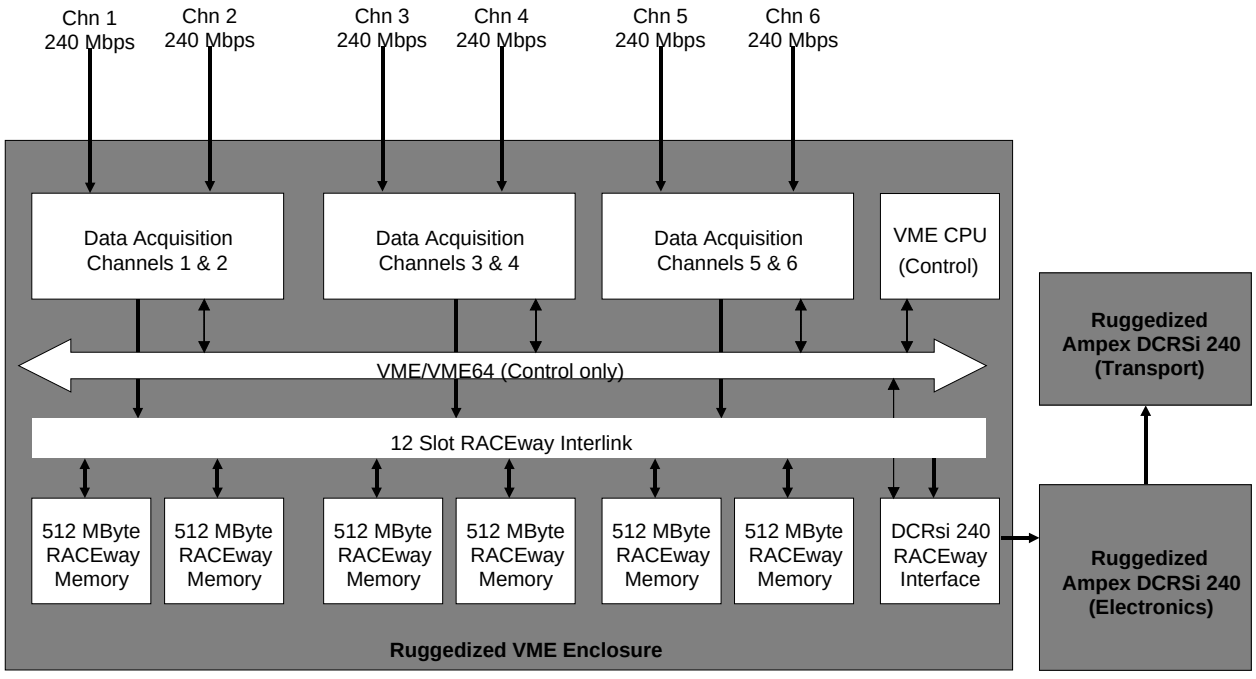

Multi-channel Analog Data Acquisition System

Currently available RACEway products listed in Table 1 can be used to construct a high-performance, multi-channel analog data acquisition and processing system (Figure 6). RACEway-based high-performance analog-to-digital conversion boards can be mixed with moderate-performance VME64 analog to digital conversion boards storing multiple streams of mixed-rate digitized data in VME/RACEway memory. A RACEway Fibre Channel board can be used to store the multi-channel data on Fibre Channel disks (or disk arrays) at rates up to 90 Mbytes/second. Data acquired in such a manner can be accessible to multiple workstations on the same Fibre Channel Arbitrated Loop (FC AL). High-rate recorder interfaces can be used to archive data. Standard VME boards can be included to implement functions such as IRIG timecode generation.

Figure 6 - Multi-Channel Analog Data Acquisition System

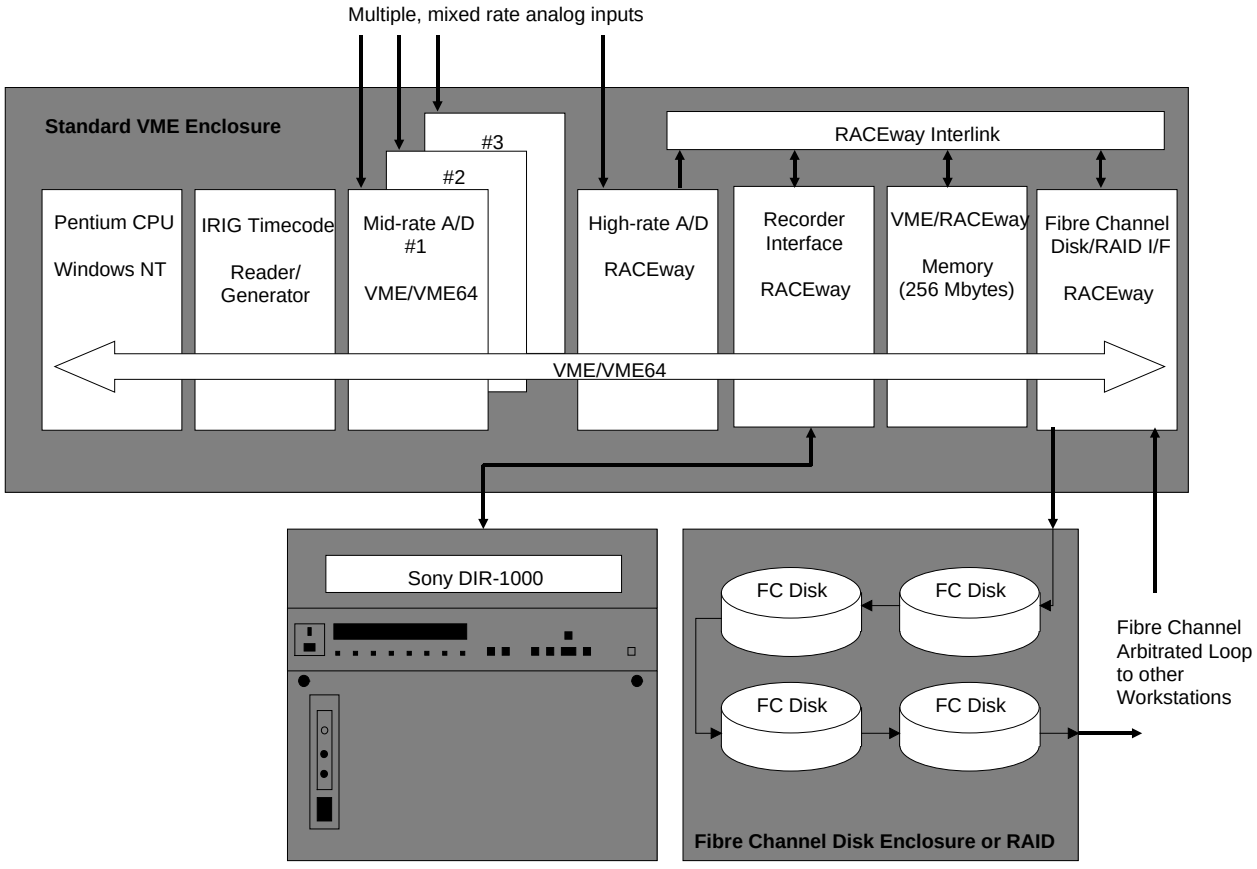

Benchmark Measurements

Benchmark testing was performed by Myriad Logic in order to determine realistic performance for various RACEway configurations. The first configuration tested consisted of two standard VME chassis, each containing one Myriad HIPPI-830R and one Micro Memory MMI 6490D. In the test the HIPPI board in chassis 1 repeatedly read fixed blocks from the memory board via RACEway and transmitted the blocks to the second chassis via the HIPPI channel. The HIPPI board in the second chassis received the blocks, storing them in the memory board via RACEway. Throughput was calculated by dividing the total number of bytes transferred over a large number of blocks (therefore, all software overhead is included in the calculation). The results measured 91.9 Mbytes/second with 1 Mbyte blocks and 95 Mbytes/second with 4 Mbyte blocks (Figure 7 graphs block size vs. measured throughput). Since the HIPPI is a 100 Mbyte/second synchronous channel, achieving 95 Mbytes/second was considered very good.

A second test was conducted with two HIPPI boards in a single chassis with a single RACEway memory board. The first HIPPI board read blocks of data from the memory board via RACEway and transmitted the data out the HIPPI channel. The second HIPPI board received the data and wrote the blocks to another area in the same memory board. This test measures RACEway throughput to a single RACEway slave board from two competing RACEway master boards and thus provides a measure of the Interlink's efficiency in arbitrating requests for the same RACEway slave. In this test, performance grew with block size, rising to a maximum of 150 Mbytes/second with block sizes of 4 Mbytes. Figure 8 provides a graph of these data. This test illustrates that it is possible to sustain data rates up to 150 Mbytes/second to a single RACEway board and indicates that with multiple transactions, system level throughput can be considerably higher.

Figure 7 - HIPPI/RACEway Benchmark

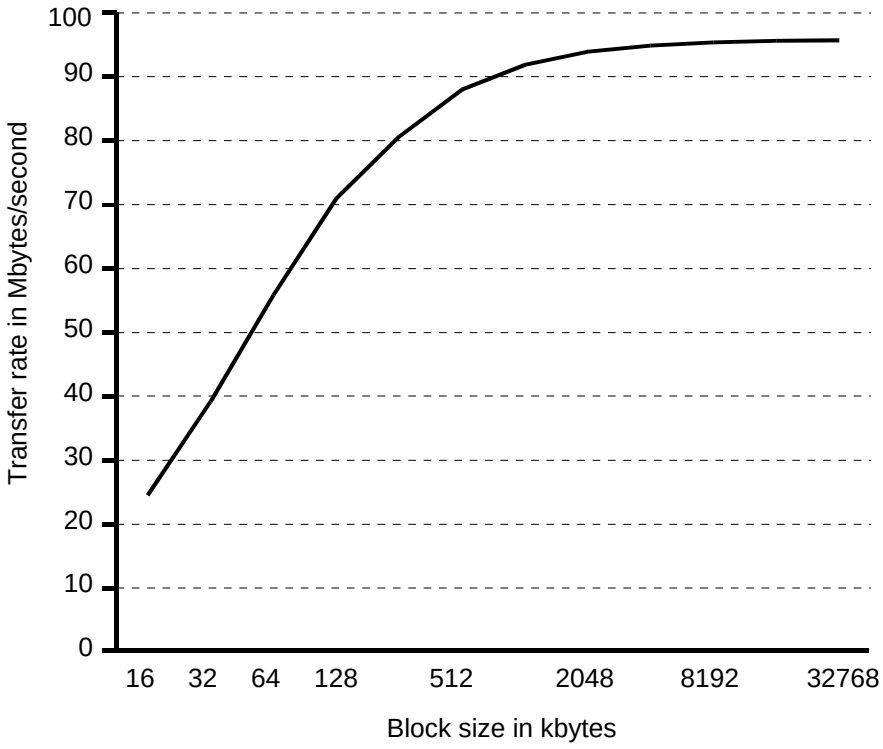

Figure 8 - Two RACEway Boards writing to a Single Memory Board

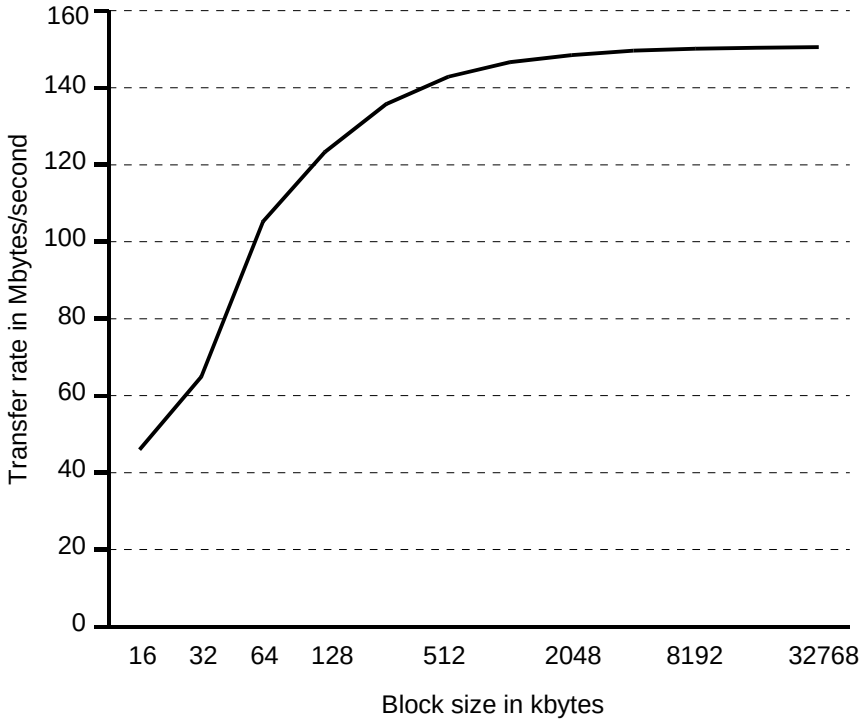

Summary

Several features of RACEway make it an excellent architecture for high-performance, real-time data acquisition, including:

1. Throughput. Since internal bandwidth is scalable, it naturally increases with the number of functional elements in the system. This is a large advantage over bus systems in which internal bandwidth decreases with the number of functional elements.

2. Deterministic Performance. The system designer can allocate separate data paths for critical real-time functions, ensuring data overruns will not occur.

3. Availability of Products. The growing number of RACEway vendors provides a wide range of off-the-shelf, high-performance products.

4. Compatibility with VME. Standard VME products -- including ruggedized enclosures, standard single-board computers, and auxiliary function equipment -- integrate easily with RACEway.

5. Preservation of Corporate Investment. RACEway allows re-use of specialized VME equipment as well as capitalizes on corporate experience with VME.

Contact Information

ANSI /VITA 5-1994, American National Standard for RACEway Interlink, July 31, 1995
VITA, 10229 N. Road Suite B, Scottsdale, AZ 85253
Phone: (602) 951-8866, Fax (602) 951-0720

Apcom, Inc., 8-4 Metropolitan Court, Gaithersburg, MD 20878-4013
Phone: (301) 948-5900, Fax: (301) 948-1631, Steven Neubauer

Atlantic Aerospace Electronics Corporation, 470 Totten Road, Waltham, MA 02154
Phone: (617) 890-4200, Fax: (617) 890-0224, Email: egraham@aaec.com, Ed Graham

Catalina Research, Inc., 1321 Aeroplaza Drive, Colorado Springs, CO 80166-2247
Phone: (719) 637-0880, Fax: (719) 637-3839, Email: cri@usa.net, Mike Bonato

Celerity Systems, Inc, 4000 Moorpark Avenue, Suite 201, San Jose, CA 95117
Phone: (408) 247-3383, Fax: (408) 247-3381, Email: mikekw@aol.com, Mike Wojnar

Cypress Semiconductor, 3901 North First Street, San Jose, CA 95134
Phone: (408) 943-2600, Fax: (408) 943-6859, Email: drh@cypress.com, Dave Horton

Echotek Corporation, 555 Sparkman Drive, Suite 400, Huntsville, AL 35816
Phone: (205) 721-1911, Fax: (205) 721-9266, Email: echotek@traveller.com, Gary Turchetta

MicroMemory Inc., 9540 Vassar Avenue, Chatsworth, CA 91311
Phone: (818) 998-0070, Fax: (818) 998-4459, Email:sales@umem.com, Mose Jadon

Mercury Computer Systems, Inc., 199 Riverneck Road, Chelmsford, MA
Phone: (508) 256-1300, Fax: (508) 256-5205, Email: mcleod@mc.com, Leigh McLeod

Mizar Corporation, 2410 Luna Road, Carrollton, TX 751006
Phone: (972) 277-4600, Fax: (972) 277-4666, Email: info@mizar.com, Keith Burgess

Myriad Logic, Inc., 1109 Spring Street, Silver Spring, MD 20910
Phone: (301) 588-1900, Fax: (301) 588-0605, Email: dbenedict@myriadlogic.com, Debra Benedict

Pentland Systems, Ltd., 1212 Baxter Drive, Plano, TX 75025
Phone: (214) 517-9343, Fax (214) 517-0070, Email: pentland@aol.com, Tracy Dorset

VMetro, Inc., 1880 Dairy Ashford #535, Houston, TX 77077
Phone: (713) 584-0728, Fax: (713) 584-9034, Email: sgonzalez@vmetro.com, Stacey Gonzalez