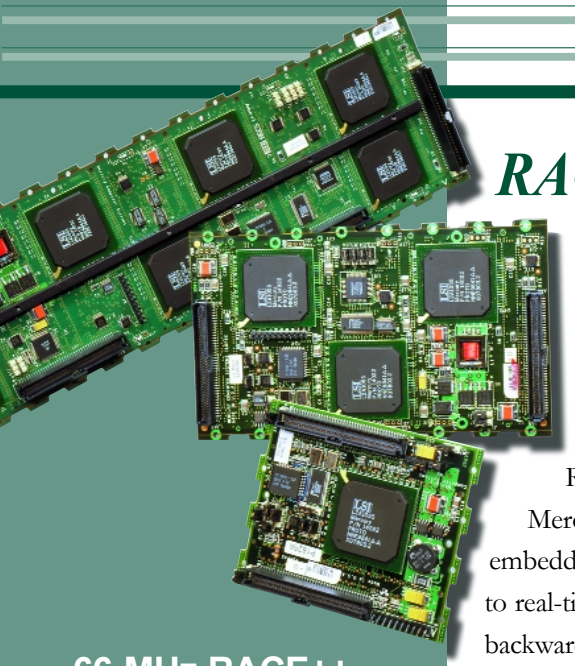

Computer Systems, Inc.
MERCURY

RACE++ Series

RACEway Interlink Modules

66-MHz RACE++ Switched Interconnect

Adaptive Routing

More than 2.1 GB/s Bisection Bandwidth

Deterministically Low Latency

Low Power, Field-Tested Design

Flexible Configuration Up to 20 Slots

RACE 1.0 Compatibility Mode

The RACE++™ Series RACEway Interlink modules from Mercury Computer Systems bring embedded high-performance computing to real-time, VME-based systems. As a backward-compatible upgrade, RACEway Interlink modules transform the topology of an existing VMEbus chassis from a single transaction bus to a scalable, real-time communication fabric capable of more than 2 GB/s of bisection bandwidth.

RACEway Interlink modules add interboard bandwidth to VME-based systems by providing multiple, simultaneous, high-speed communication paths between VME boards interfaced to the RACEway Interlink standard. In addition to increased bandwidth, the RACEway Interlink offers low latency and priority control, which are essential to applications fed by real-time sensor I/O.

Open and Standard Interface

The RACEway Interlink ANSI/VITA standard provides system designers with the option to interface value-added technologies, such as application-specific boards that contain proprietary hardware designs, with RACEway Interlink modules.

Correspondingly, the RACEway Interlink Standard has received wide acceptance with defense contractors, commercial OEMs, and VME board manufacturers. To date, there are dozens of companies that have implemented nearly 100 different boards with a RACEway interface.

VME Chassis with Interlink

Mechanically, the RACEway Interlink modules mount on the backplane of a VME chassis similar to VSB backplane modules. Electrically, these modules are connected to the VME slots through the P2/J2 chassis backplane connector. Interlink modules are available in both commercial and ruggedized VME versions.

The RACE Architecture

High-bandwidth interboard communication by itself is not sufficient to provide high-performance computing in embedded systems. The RACE™ architecture provides an open framework for evolving balanced computing systems in which different types of compute nodes (CNs), I/O nodes, and the topology can be tailored for a given application.

The RACE architecture consists of RACEway Interlink modules, compute nodes, memory-only nodes, RACEway crossbars, special-purpose devices, and I/O nodes from both Mercury and third parties (see Figure 1). All can be interconnected by the RACEway crossbar fabric to create optimized configurations for the needs of a specific application.

Interoperability between RACEway and other interconnect architectures is well established, with direct bridges to VME, PCI, Fibre Channel, HiPPI, Front Panel Data Port (FPDP), and others. Some of these bridges are Mercury's design and some are third-party designs.

RACEway Interconnect Fabric

The RACE++ Series RACEway Interlink modules connect multiple RACE++ and RACE 1.0 VME boards. These modules are composed of the same crossbar ASICs that are used to interconnect the individual nodes on a RACE++ Series VME board.

Crossbars on VME boards connect multiple nodes within the boards. Together, the RACEway crossbars on the Interlink modules and on the VME boards make up the RACEway interconnect fabric. The RACEway fabric provides a uniform inter-node communication medium that connects processors, I/O devices, and standard bus interfaces, such as VME and PCI, in a consistent way throughout the system. This simplifies software and allows users to easily scale the number of processing elements both within a board and between boards for their application.

The RACE++ crossbar is an 8-port crossbar implemented in a single low-power ASIC. The crossbar can perform four simultaneous transfers of 267 MB/s for a total of 1.07 GB/s, and can broadcast to seven ports at an aggregate rate of 1.87 GB/s.

In addition to providing scalable, high throughput, the RACEway interconnect fabric is also designed to meet low-latency, real-time data transfer requirements. When making a connection through a RACE++ system at 66.66 MHz, each crossbar along the selected path adds only 75 ns to the latency. Even in the face of contention for a data path, the RACEway interconnect fabric provides deterministically low latency through a crossbar design that provides programmable priorities under application control. For example, priority assignment allows an I/O device to preempt other data transfers in order to distribute real-time sensor data throughout the system in a deterministic manner.

Figure 1 Sample RACEway System Configuration

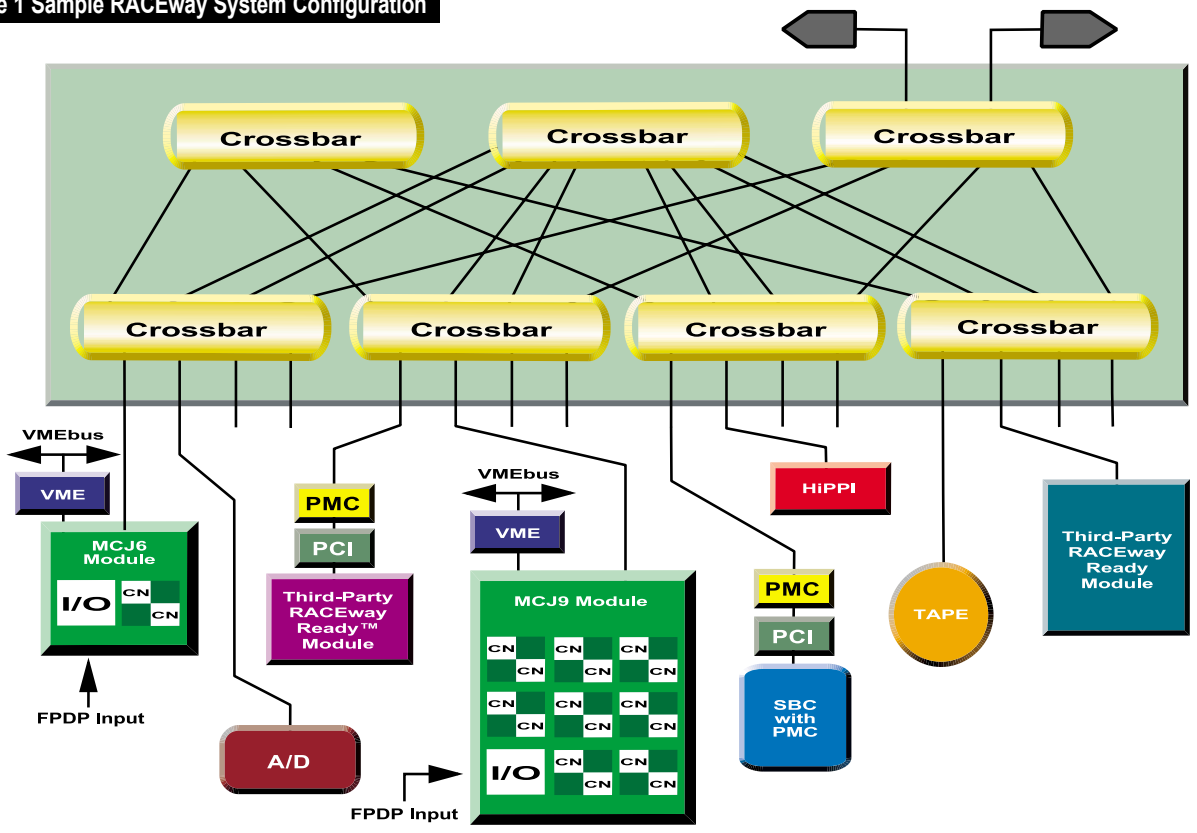

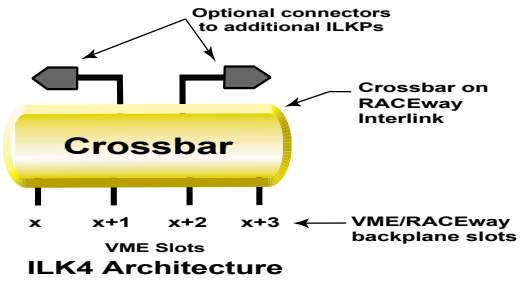

Figure 2 ILK4P Architecture

Architecture of Interlink Modules

The ILK4P Interlink module consists of a single RACEway crossbar that spans four VME slots and has two connectors for expansion to other Interlink modules. Figure 2 shows the architecture of the ILK4P. The ILK4P is particularly adapted to low-end systems that require up to four boards.

ILK8P, ILK12P, and ILK16P

The ILK8P, ILK12P, and ILK16P Interlink modules are designed for larger VME multicomputer applications. They provide a high-performance interconnection among 8, 12, and 16 slots, respectively. They each have two connectors to allow for expansion. The architecture of these ILKPs is given in Figures 3 through 5.

These devices have network topologies that provide a high interconnection bandwidth by exploiting parallel RACEway paths. The parallel paths on the Interlink module, which support a given set of concurrent data transfers, are automatically selected by the RACEway crossbar ASICs. If one path is already busy with a transfer, the free path is automatically selected by the crossbar logic. This is called "adaptive routing" and is described in detail in the next section.

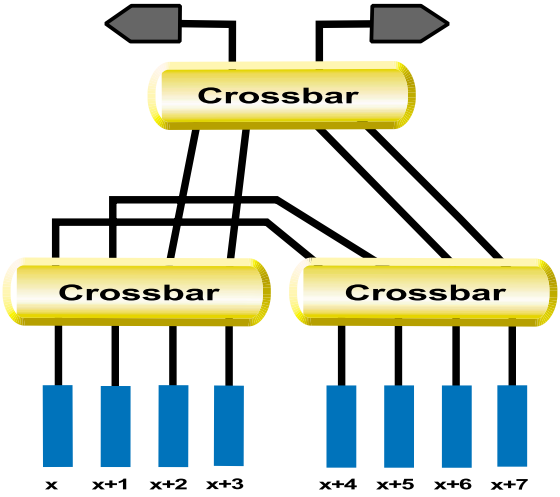

Figure 3 ILK8P Architecture

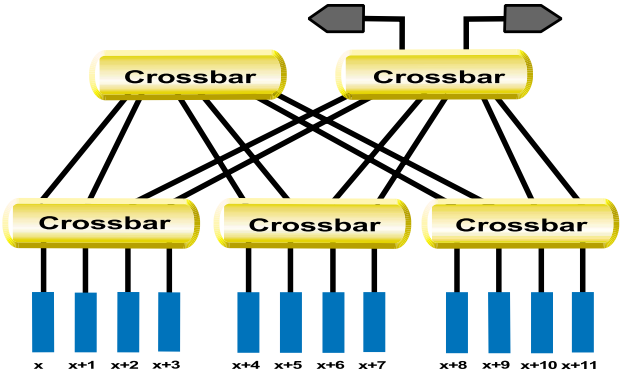

Figure 4 ILK12P Architecture

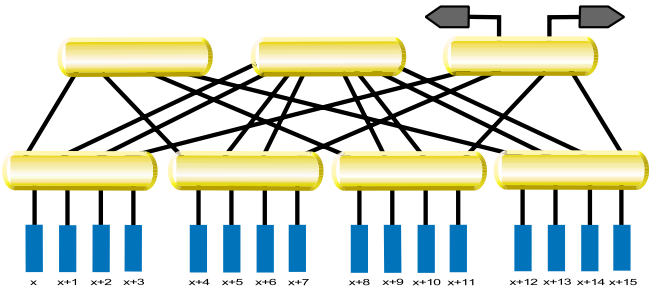

Figure 5 ILK16P Architecture

Adaptive Routing

Adaptive routing (also referred to as auto route path selection) is a remarkable RACEway feature that greatly increases the utilization of the network resource by automatically routing traffic toward underutilized links. Figure 6 illustrates how the RACEway crossbar ASIC implements adaptive routing. Figure 6a shows a transfer in progress between CN1 and CN3. Figure 6b shows how the crossbar ASIC automatically chooses the free available path through the Interlink crossbar network between CN2 and CN4 for a second transfer. Primarily a performance feature, adaptive routing is also an ease-of-use feature that frees the programmer from the details of path routing. Distributed matrix transposes, or "corner-turns" such as those that

occur in SAR applications, achieve very high interprocessor communication bandwidth with this feature. Adaptive routing is the default routing option with the ILK8P, ILK12P, and ILK16P devices.

Path Priorities

Important transfers such as sensor I/O or synchronization messages can be programmed to preempt lower-priority transfers, ensuring deterministic message transfer time regardless of lower-priority internode traffic. Figure 6c shows how transfers #1 and #2 (in Figure 6b) are preempted when a higher-priority transfer #3 takes place from CN5 to CN1. The lower-priority transfer #2 from CN4 to CN2 can resume immediately by adaptively taking a different route. The lower-priority transfer #1 will resume when transfer #3 completes.

Figure 6a-6c

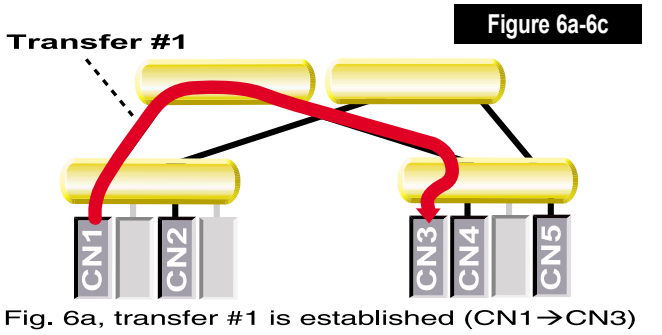

Fig. 6a, transfer #1 is established (CN1→CN3)

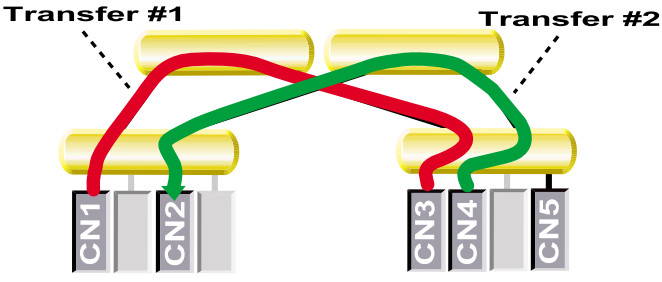

Fig. 6b, transfer #2 (CN4→ CN2) automatically routes to free path

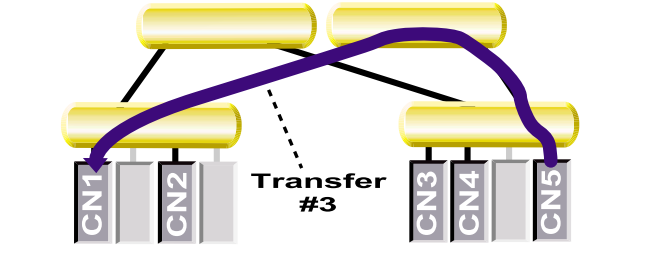

Fig. 6c, a high-priority transfer (#3) preempts other transfers (any route can be chosen)

ILK1P Module

The ILK1P is a single VME slot Interlink module. It consists of a one-slot connector containing clocking circuitry and a ribbon cable. The clock frequency of the ILKP1P is fixed at either 40 MHz or 66.66 MHz, depending on the model ordered.

The ILK1P gives the system designer additional configuration flexibility in two ways. First, ILK1P devices can be used to expand the ILK4P, ILK8P, ILK12P, and ILK16P Interlink modules (see configuration rules for details). Second, two ILK1P modules can be connected with an Interlink cable to create a low-cost connection between two RACEway Interlink-compatible VME boards.

Aggregate and Bisection Bandwidth

The peak aggregate performance of Interlink modules is defined as the total available bandwidth of the interconnection system on the Interlink module. This is the sum of the maximum number of simultaneous inter-board transfers.

When a network is divided into two equal halves in such a way that minimizes the bandwidth of the interconnection between the two halves, the resulting bandwidth between these two halves is defined as the bisection bandwidth. For the ILKPs, the peak aggregate and peak bisection bandwidths are equal and are given in Table 1.

Table 1: Bisection Bandwidth of ILKPs

ILKP Type	Number of RACEway Links Crossing Bisection	Peak Bisection* Bandwidth (GB/s)	
		40-MHz Clock	66.66-MHz Clock
ILK4P	2	.32	.53
ILK8P	4	.64	1.07
ILK12P	6	.96	1.60
ILK16P	8	1.28	2.13

* For each ILKP, the aggregate and bisection bandwidths are equal.

Physical Description

RACEway Interlink modules attach to the user-defined pins on P2 "wire wrap posts" of a VME backplane. The backplane can have either 3-row connectors (96-pin) or 5-row connectors (160-pin) for P2. The chassis/backplane should be ordered with these wire wrap posts, also called VSB pins, in place (electrically, only rows A, C, and power/ground from row B are used for ILKP modules).

RACEway Interlink modules are attached to the VME P2 connectors in the chassis by a two-level mechanical connection involving two printed circuit boards (PCBs), as illustrated in Figure 7. The two levels eliminate the need for a high-insertion-force connection to the P2 wire posts when mounting large Interlink modules. All Interlink modules thus consist of two or more boards. The first of the boards is the interposer. The interposer has high-insertion-force connectors spanning one or four VME slots. These plug onto the P2 connector wire posts on the backplane of the chassis. The other side of the interposer contains low-insertion-force connectors that mate with the board that has the crossbars.

Top View
VME Backplane

Figure 7

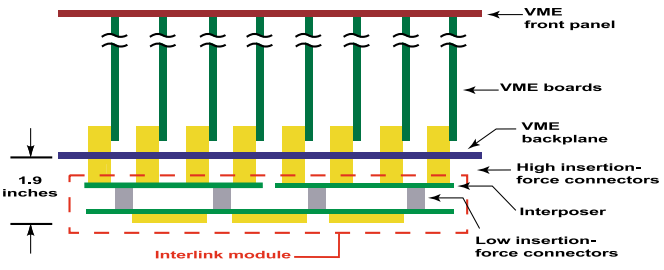

The ILK8P, for example, has two interposers, each spanning four slots. A single 8-slot board with crossbars plugs into the two interposers. The complete assembly is low profile and requires 1.9 inches of clearance from the chassis backplane. An additional 0.6 inches is required for natural convection of air for cooling.

Interlink Configuration Rules

1. Interlink modules require that the VME chassis contain standard "wire wrap posts," with shrouds, on the P2 connector of the backplane.
2. There is a different part number for the Interlink module, depending on whether the backplane has 3-row or 5-row connectors.
3. Appropriate clearance must be provided between the VME backplane and the back panel of the chassis enclosure (see Mechanical Specifications).
4. Not all RACEway slots spanned by an Interlink module need to be filled with boards (i.e., slots can be skipped).
5. VME boards inserted in RACEway slots (spanned by an Interlink) must have a valid RACEway interface or have rows A and C of the P2 connector unused.
6. A single Interlink module may not span a "split VME backplane" (two electrically isolated backplanes in the same chassis).

Additional configuration rules that apply to configurations with multiple Interlinks:

7. The number of VME chassis slots covered by an ILK4P can be expanded to one or both of the slots adjacent to the ILK4P, by adding up to two ILK1P modules (no skipped slots).
8. The number of VME chassis slots covered by an ILK8P, ILK12P, or ILK16P can be expanded by one by adding an ILK1P in a slot adjacent to the ILKP being expanded (no skipped slots). The number of VME chassis slots covered by an ILK16P can be expanded by four by adding an ILK4P in a slot adjacent to the ILKP16.
9. Interlink expansion cables (and single Interlink modules per rule 6) may not span separate VME backplanes (i.e., connection of Interlink cables across split VME backplanes or between chassis is not supported).
10. Only short Interlink cables (ILKC-A) may be used.
11. Two ILK1P modules may be connected together.

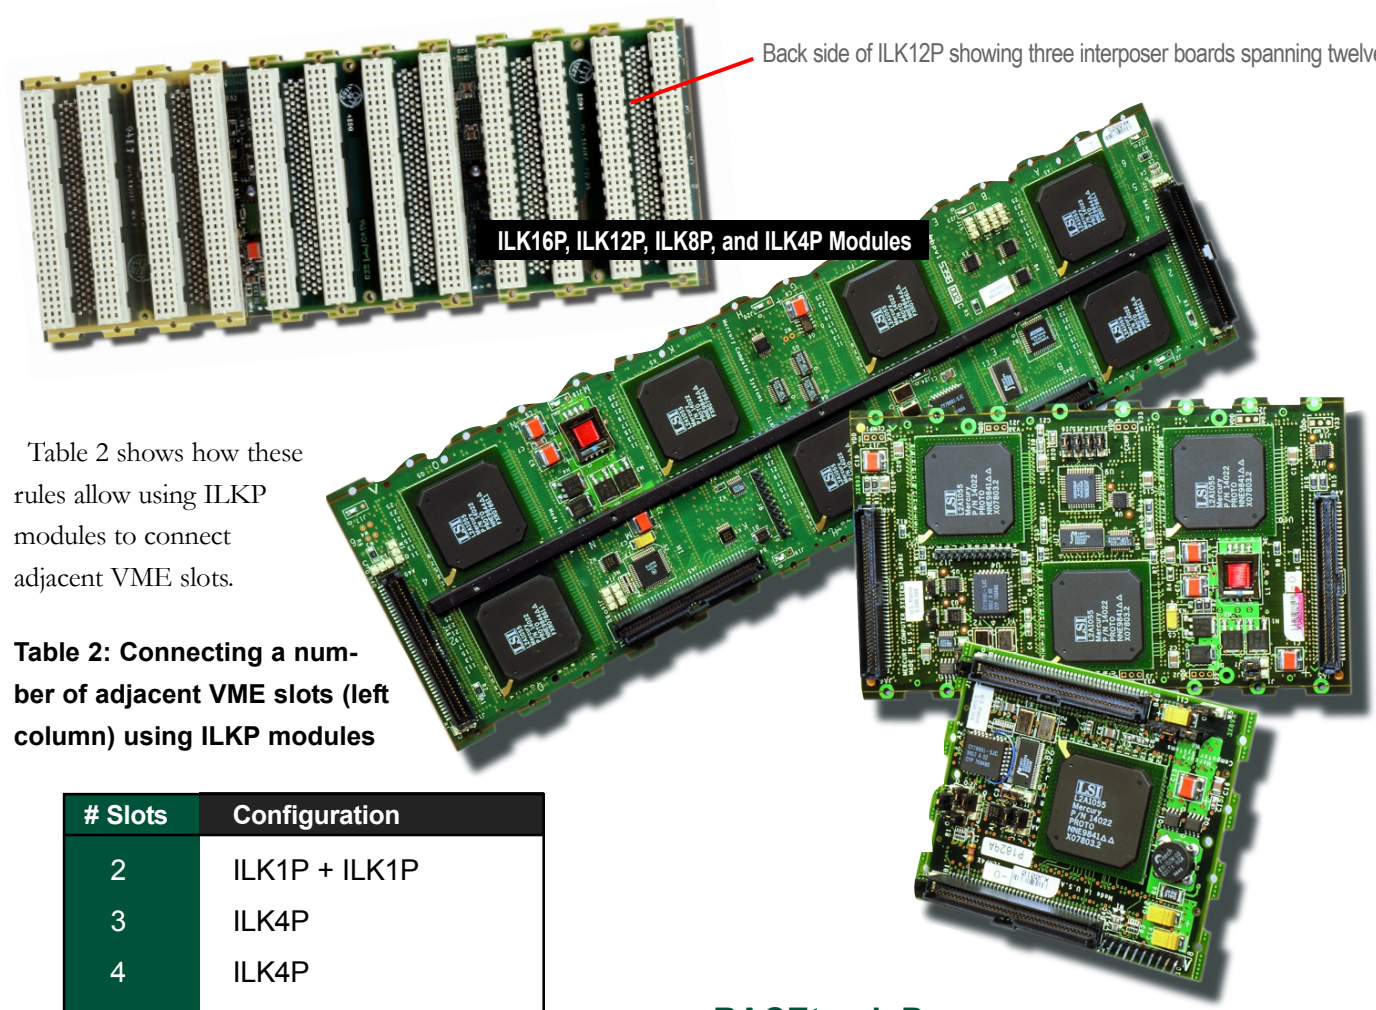

Table 2 shows how these rules allow using ILKP modules to connect adjacent VME slots.

Table 2: Connecting a number of adjacent VME slots (left column) using ILKP modules

# Slots	Configuration
2	ILK1P + ILK1P
3	ILK4P
4	ILK4P
5	ILK1P + ILK4P
6	ILK1P + ILK4P + ILK1P
7	ILK8P
8	ILK8P
9	ILK1P + ILK8P
10	ILK12P
11	ILK12P
12	ILK12P
13	ILK1P + ILK12P
14	ILK16P
15	ILK16P
16	ILK16P
17	ILK1P + ILK16P
18	ILK4P + ILK16P
19	ILK4P + ILK16P
20	ILK4P + ILK16P

RACEtrack Program

Mercury offers a service to assist with the design of interfaces to the RACEway interconnect fabric. The RACEtrack™ program is a package of consulting services and documentation for users who design boards that interface to the RACEway interconnect. Available for purchase through the Mercury Alliance Program, the RACEtrack program includes a RACE I/O software driver, diagnostics, examples, schematics, and a selection of prepaid support and services. Through this program, third parties and customers can build their own modules, I/O devices, or special-purpose boards that interface to the RACEway Interlink fabric.

RACE++ Technology

As the latest version of the RACE architecture, RACE++ technology represents an architectural evolution that includes increased communication speed, more richly connected topologies, and augmented adaptive routing. Together these enhancements amount to a performance revolution yielding significantly higher bisection bandwidth and lower latency to attack the most challenging real-time problems.

The RACE architecture includes motherboards and RACEway Interlink modules that connect motherboards. These RACEway Interlink modules are Mercury's implementation of the RACEway Interlink standard, which was introduced in 1993 and adopted as an ANSI/VITA standard (ANSI/VITA 5-1994) in 1995. The RACE++ version of the Interlink modules are referred to as ILKPs. The RACEway clock generated by the ILKPs can be set at either 40 or 66.66 MHz, using field-settable jumpers.

RACE++ technology builds on the strengths of the original RACE architecture (RACE 1.0). It implements evolutionary enhancements, evenly distributed over the key features of the RACE 1.0 crossbar: connectivity, speed, and adaptive routing. This balanced set of enhancements yields impressive gains in overall system interconnect performance. RACE++ technology delivers more powerful systems with greater bandwidth in an operational environment fully suited for long-term reliability in harsh conditions.

The RACE++ interconnect implements a superset of the 1995 RACEway standard. It allows system designers to build computers with more complex internal data network topologies. This allows larger, faster, more powerful multicomputer systems. Compared with the 1995 RACEway standard, the gain in overall system-wide sustained bandwidth can reach factors in excess of three.

This achievement is made possible by strategic, manageable increases in the performance of the RACE crossbar ASIC. The crossbar is the cornerstone of the RACE architecture and provides connectivity between nodes within a RACE system.

The table below shows the key crossbar changes between the RACE 1.0 and RACE++ interconnect, and their impact on overall system performance.

The RACE++ fabric is constructed from 8-port crossbars. The resulting network topologies and improved adaptive routing provided by the ILKPs offer substantial improvements in performance over their earlier ILK counterparts, even when operated at 40 MHz.

Bisection Bandwidth of ILKPs (RACE++) and ILKs (RACE 1.0)

When compared at 40 MHz, the RACE 1.0 ILK4 and ILK8 have the same bisection bandwidth as their RACE++ counterparts, the ILK4P and ILK8P. The ILK12P and ILK16P, however, both have an increased bisection bandwidth: from 4 RACEway links (640 MB/s) for the ILK12 to 6 RACEway links (960 MB/s) for the ILK12P, and from 4 RACEway links (640 MB/s) for the ILK16 to 8 RACEway links (1280 MB/s) for the ILK16P. This means that the ILK12P has 1.5x the bisection bandwidth of an ILK12, and the ILK16P has 2x the bisection bandwidth of an ILK16, when compared at 40 MHz. When operated at 66.66 MHz, the increased clock frequency gives the ILKPs another factor of 1.67x, making the bisection bandwidth of an ILK16P operating at 66.66 MHz a factor of 3.33 times larger than that of an ILK16.

In addition to the increased bisection bandwidth, ILKPs offer increased performance over their ILK counterparts through better adaptive routing. The benefits of improved adaptive routing come into play when there is contention between two or more concurrent RACEway transfers.

Key Crossbar Attribute	Crossbar Feature Improvements	Overall System Benefits
Connectivity	From 6 to 8 ports	Denser topologies, more paths
Speed	From 40 to 66.66 MHz; 40 MHz still supported for compatibility	More bandwidth while preserving signal integrity
Aggregate Transfer Bandwidth	From 480 to 1067 MB/s	Significantly increased bisection bandwidth
Adaptive Routing	From 2 to 8 adaptive routing ports	Greater network resource utilization
Low Power	UNCHANGED!	One chip, one watt = more bandwidth with no power penalty

Electrical Specifications

	ILK1P	ILK4P	ILK8P	ILK12P	ILK16P
Power Consumption (Watts)	0.5	1.75	10	13	13
Input voltage ALL	4.75 to 5.25 volts (3.3 volt power is not used from backplane)				

Mechanical Specifications

	ILK1P	ILK4P	ILK8P	ILK12P	ILK16P
Weight (pounds)	0.1	0.3	0.7	1.1	1.4
Dimensions (inches):					
Length	0.76	3.12	6.32	9.52	12.72
Width	3.74	3.74	3.74	3.74	3.74
Height	2.1	1.9	1.9	1.9	1.9

Module height from the VME backplane includes cable. Modules are shipped by default with breakaways intact. The width without breakaways is 3.53 inches.

Environmental Specifications

Mercury offers multiple categories of environmental ruggedness for the RACE++ Series RACEway Interlink modules. The specifications detailed below are for the RACE++ Series RACEway Interlink modules for standard commercial environments.

Operating temperature	0°C to 50°C up to an altitude of 10,000 ft
Storage temperature	-40°C to 85°C
Relative humidity	10 to 90% (non-condensing)

Ordering Information

Part Number	Model	Description
910-08034	ILK1P	Single-slot RACE++ Interlink (5 row)
910-08009	ILK4P	4-slot RACE++ Interlink (5-row)
910-08010	ILK8P	8-slot RACE++ Interlink (5-row)
910-08011	ILK12P	12-slot RACE++ Interlink (5-row)
910-08012	ILK16P	16-slot RACE++ Interlink (5-row)

Call Mercury for information on ordering 3-row versions of all ILKP modules.

mc.com Find us on the web at www.mc.com

Computer Systems, Inc.
MERCURY
The Ultimate Performance Machine

199 Riverneck Road
Chelmsford, MA 01824-2820 U.S.A.
978-256-1300 · Fax 978-256-3599
800-229-2006 · http://www.mc.com
NASDAQ: MRCY

Nihon Mercury Computer Systems K.K.
No. 2 Gotanda Fujikoshi Bldg. 4F
5-23-1 Higashi Gotanda
Shinagawa-ku, Tokyo 141-0022
JAPAN
+ 81 33 473 0140
Fax + 81 33 473 0141

Mercury Computer Systems, SARL
Immeuble le Montreal
19 bis, avenue du Quebec
Villebon
91951 COURTABOEUF CEDEX
FRANCE
+ 33 (0) 1 69 59 94 00
Fax + 33 (0) 69 86 10 87

Mercury Computer Systems, Ltd.
Campbell Court, Unit 19
Bramley, Tadley
HANTS RG26 5EG
UNITED KINGDOM
+ 44 1 256 880090
Fax + 44 1 256 884004