By Remy David,
Director of Marketing and Technology,
GESPAC S.A.

Why CompactPCI will replace VME

In just a year, the CompactPCI bus has risen to prominence in the board industry worldwide, prompting users to compare it with VME and to wonder if it could replace VME in many applications. There are two ways to compare buses. The first is to look at the intricate details of the bus signals, bus transceiver technology and timing diagrams. From this point of view, VME (and VME64x) fares as an outstanding technology with plenty of growth potential. The other way is to see what silicon components, system and software facilities a given architecture offers to simplify system design. From this perspective, PCI and CompactPCI offers unique possibilities thanks mostly to the broad silicon support and "system-oriented" concepts like Plug and Play. While the installed base of applications and customers ensures VME many more years of existence, these powerful force are already at work to make CompactPCI the dominant industrial OEM bus at the end of the decade.

Historically, new industrial bus star
ly have originated from a numb
As semiconductor companies in
microprocessors, they also bui
board and system products designed r
plify the use and speed the acceptance

any given chip or computer maker.

An active, but commercially less successful source of buses has been institutions like the IEEE. Notable ventures include the STE bus and more recently Futurebus. While none of these buses was ever endorsed by a large company or gained market acceptance, they have contributed to the development of new technologies that eventually made their way in many of the buses that are in use today.

It is fair to say, however, that the most important contribution to the bus/board market have been done by the Computer and the Semiconductor industries.

BOARD LEVEL VS. SYSTEM LEVEL BUSES

Buses in system hierarchy

Buses are elaborate technologies and specialists have dedicated careers to refine them. They can also be emotional subjects; as recent "battle of the buses" have shown. But one must not lose sight of the big picture. Buses are essentially communication schemes designed to build systems. Figure 1 shows a layered

Difference between computer and semiconductor maker's buses.

Buses developed by chipmakers, like VME, tend to focus on the elementary functions of the system. In the

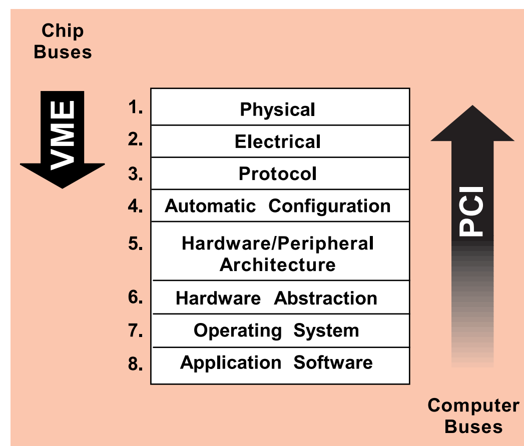

Figure 1. Layered model of the elements of a typical open system

VME systems are gaining (not losing) interest!

- Many companies are defecting from CPCI back to VME
- CPCI hype is being recognized for what it is

Global Merchant Embedded Slot Card Market, 2002 & 2003

US$ in Millions

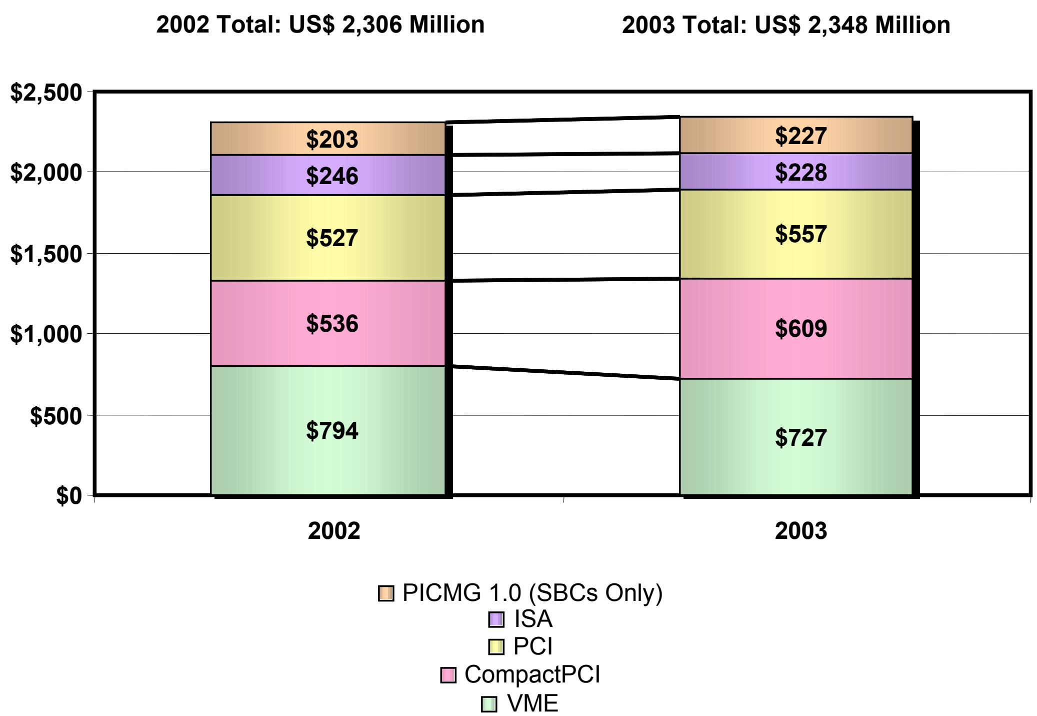

By Bus Architecture 2004 (by $ volume)

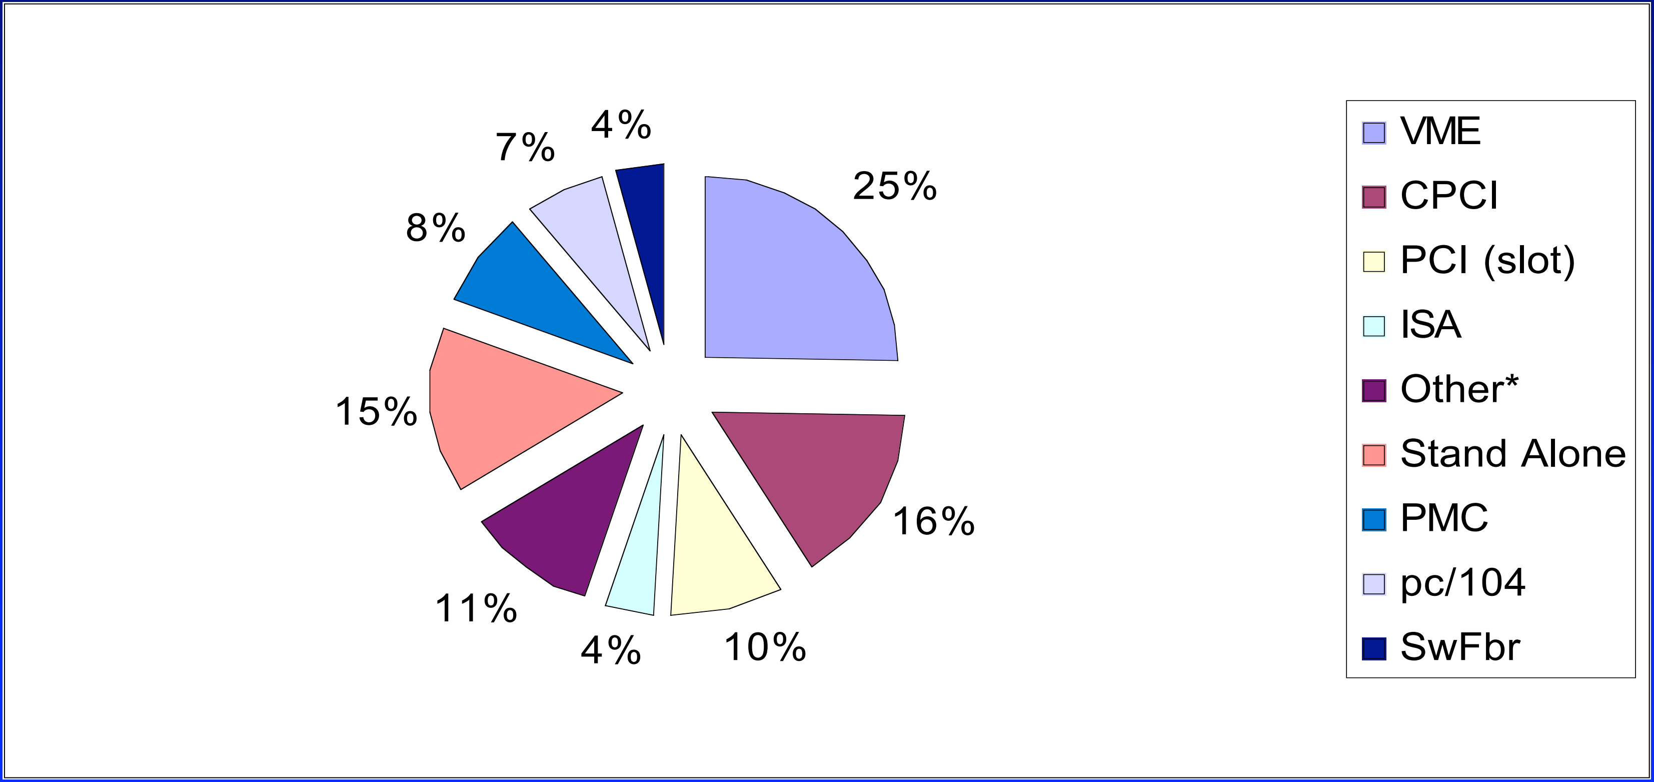

Current and Forecast Markets

- **VME: $1.15 billion – 5% - 10% continued growth**
  - Strong mil-COTS segment
  - Realtime apps in wireline, wireless & 3-G, Industrial Automation
- **cPCI: $200 million including OS and firmware**
  - Soft market – packet switching and "last mile" apps in doubt
  - Custom board markets
- **PC 104/+: $50 million**
  - Niche market
  - Threat from FPGAs
- **ATCA: cure looking for a disease**
  - Strong possibilities for high bandwidth (e.g., WDM) apps
  - 3-G and wireless possibilities
  - Solution for a world of connected devices?
  - Multi $billions if it takes hold – middle ground not likely

Jerry Krasner, Ph.D., MBA

- VME vs. cPCI vs. PCI: We know the players - but on whose court and by what rules are we keeping score?

- The High End Board Market will be dominated by VME and the 5-row DIN Connector
- cPCI and the 2mm HM connector will dominate Telecom and the 3U Industrial and Transportation markets

VME

- will continue to dominate multiprocessing applications
- is driven by long-life cycle applications
- COTS will be a major driving force
- Presidential politics will drive military design starts
- higher volume customers will switch to contract manufacturing
- VME's incorporation of high-speed serial data flow will extend its high-end market dominance

cPCI

- number of legal backplanes under PICMG spec: Asset & Problem
- choice of telecom engineers for software compatibility, rear-connect phone lines allowing unrestricted card insertion and front end card swapping.
- 6U form factor provides additional pin-outs
- Long lead-time design process by telco engineers can be to cPCI's Advantage.
- In House or Contract Manufacturing Decisions Could Hamper Merchant Opportunities
- need for hot swap: technical and patent issues

Why cPCI?

- All present day computers have some sort of PCI bus (including VME SBCs)
- Only 47 lines for 32 bit bus
- Inexpensive silicon & connectors
- IP (Xilinx, opencores.org, etc.)
- VME interface chips (e.g. Tundra Tsi148) are bridges to PCI
- PMC cards used in VME modules
- “Slave” terminated block read

PICMG

The PCI Industrial Computer Manufacturer's Group

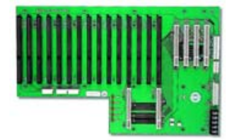

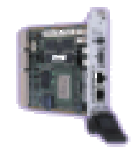

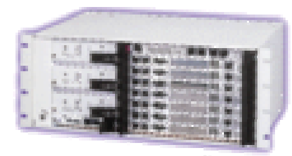

- A consortium that has standardized ISA and PCI technologies for industrial backplane applications
- Founded in 1994 with some 600 member companies today
- Have authored a broad family of specifications that define complete hardware architectures
  - Physical board sizes
  - Power, thermals, and electricals
  - System management
  - Backplane layouts and protocols

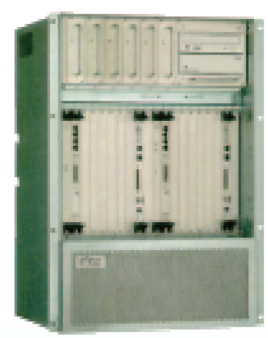

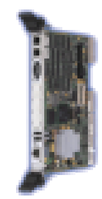

XILINX

From cPCI FAQ:

Q: What other functions could benefit from CompactPCI?

A: CompactPCI can benefit all applications requiring very high data transfer rates. Data communication interfaces such as ATM and broadband ISDN are good examples. In the [illegible] fast multi-channel data acquisition cards could also benefit from CompactPCI. Many of the most exciting applications are probably yet to be invented, but if history is any indication, the sophistication of systems will increase to use all available computing bandwidth that CompactPCI computers have to offer.

Q: What are the software implications of PCI and CompactPCI?

A: The PCI architecture, developed by Intel, has been carefully planned to simplify the software integration of a peripheral device. For example, all PCI or CompactPCI devices have a set of 256 registers which contain information on the device identity, as well as a great deal of software programmable parameters such as address maps, or interrupt types and levels. As a result, the system CPU can automatically detect and identify a device on the bus and configure it without the need for jumpers on the peripheral. PCI is a key element of the "Plug and Play" concept.

Back to Top

Q: What are the system implications of CompactPCI?

A: Modern computer architecture has an internal PCI bus, which usually supports PCI add-on slots. This is the case for nearly all Pentium PCs, Alpha workstations, and PowerPC systems based on the PREP or CHRP reference platform standard. CompactPCI makes it possible to build any computer compliant with these hardware system designs. As a result, CompactPCI systems can be built using standard components and can run practically any operating system and thousands of application software packages without modification.

Q: How many connectors are used on 6U CPCI plug-in boards?

A: Four or five 2-mm connectors are used on 6U CPCI plug-in boards. Identical connector configurations are used on the upper and lower 3U sections of a 6U board. The CPCI bus is on the bottom connector. The upper connector set can be used for a second CPCI bus or for 220 user I/O pins and 44 grounds. A 19-row version of the 2-mm connector can be placed between the upper and lower connector sets providing 95 more user I/O pins and 19 grounds.

Q: Does CompactPCI support all of the "Plug-and-Play" configuration features in the PCI specification?

A: Yes, including the use of PCI-to-PCI bridges. The CPCI specification details a backplane power supply slot for plug-in PSUs. The PSU can be powered from 110/220 VAC or 48 VDC, and can supply +5.0V, +3.3V, and +/-12V. Pins are also defined for "Current Share", "Supply Fail", and "Supply Derate", and "Inhibit". The "Current Share" signal is used for two or more PSUs running in a redundant configuration. The "Supply Fail" and "Supply Derate" signals indicate the viability of the PSUs to the CPCI system. The "Inhibit" signal is usually connected to a front panel switch and is used to turn off the PSU.

Q: Can you combine CompactPCI and VME64x?

A: Yes, since VME64x and CompactPCI use the same Euroboard mechanical packaging, it is possible to combine both architectures into a single chassis. A special VME64x bridge connector pinout has been proposed for VME64x to CPCI bridgeboards. This special connector has all of the VME64x signals on a single 2-mm connector. A bridgeboard would connect to the VME64x bus with the top connector and the CPCI bus with the bottom connector. This implementation would require a special backplane with both VME64x and CompactPCI on the same PCB. The same expert that simulated the design of CompactPCI validated the VME64x 2mm-bridge pinout.

Q: Are there system integration issues when combining CPCI and VME64x on the same backplane?

A: When the CPCI is the master bus it is normally placed on the left side of the chassis. The bridgeboard would be the last CPCI slot and the first VME64x slot. The means that the bridge would have to be the VME64x system controller. This would be easy to implement with the Tundra Semiconductor Universe chip.

When the VME64x is the master bus the bridgeboard is the last VME64x slot and the first CPCI slot. The bridgeboard would then have to support the interrupt capabilities normally associated with a CPCI CPU board. It would be possible to layout the CPCI bus backwards with the CPCI CPU at the far right. This would make the bridge board a slave to both the VME64x and CPCI bus.

A CODA interface to IU FADC PCI prototype exists.

Backplane Conclusions (1)

- Backplanes are the foundation of modular system designs
- Switched fabrics are the preferred architecture for modern backplane implementations
  - They provide higher performance through their support of simultaneous transactions
  - They scale well
- Serial switched fabrics are the ultimate backplane solution
  - Limited pin count
  - Reduced power
  - Robust aggregation
- Serial switched fabrics make Full Mesh topologies practical

Switched Fabrics

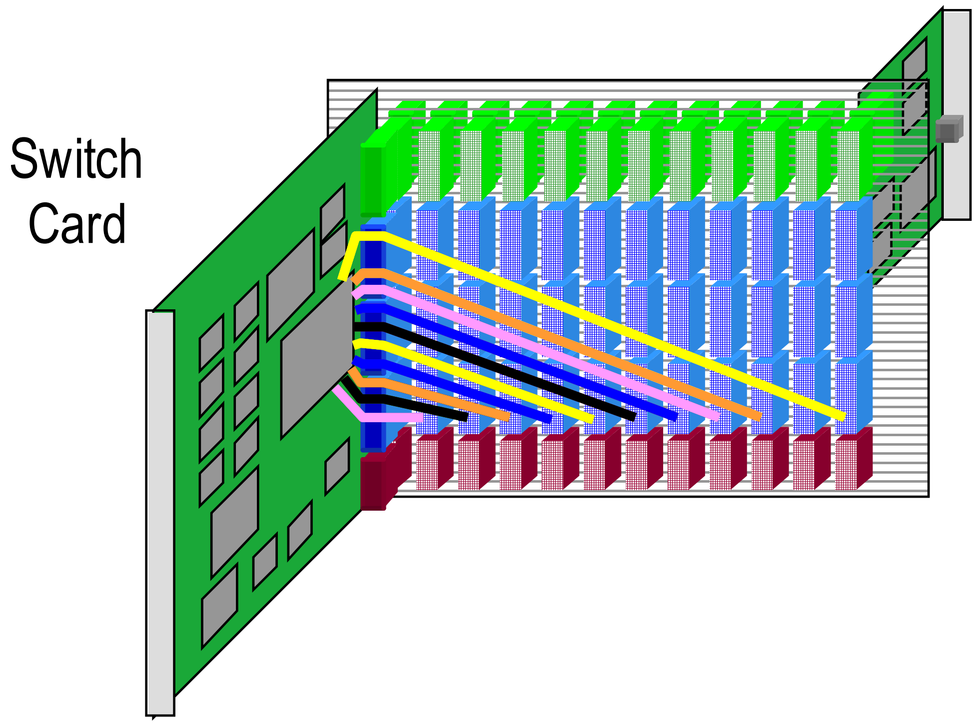

- Node cards connect to each other through a Switch card
- Switches can support simultaneous traffic between different node card pairs
  - Increasing available bandwidth

PICMG 2.16

Topologies

- The PSB links runs across a Star or Dual Star configuration (not a bus) in a centralized topology.
- Each line interconnecting a Node Slot and Fabric Slot represents a Link that is a 10/100/1000 Mbps full duplex Ethernet connection

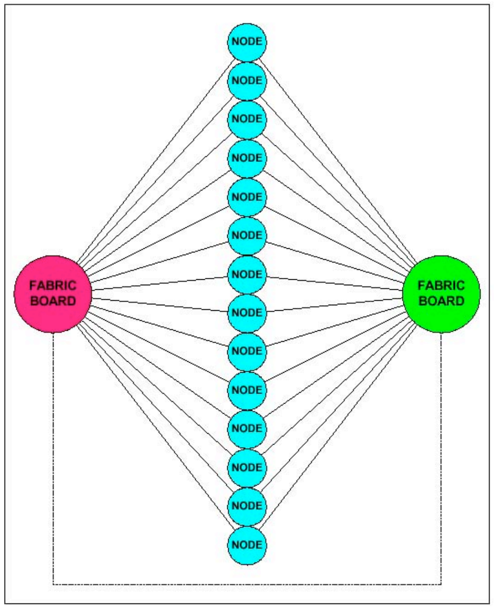

Dual Fabric PSB Topology

(Two Link Ports)

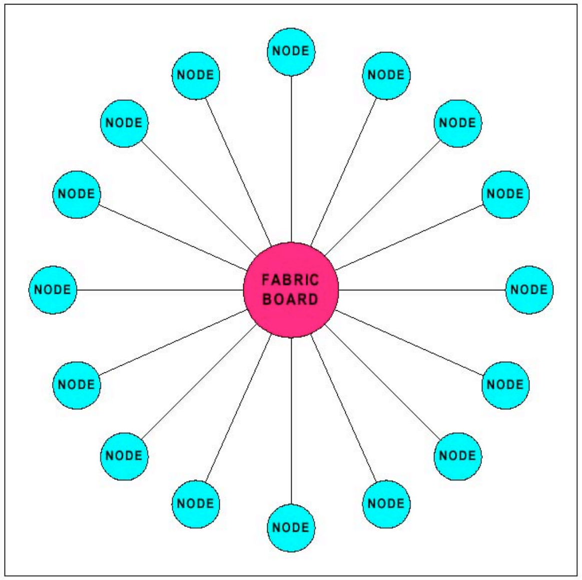

Single Fabric PSB Topology

(One Link Port)

PICMG 2.16

bus tronic
An Elma Company

About PICMG 2.16 (CompactPCI Packet Switching Backplane cPSB)

- cPCI can still be used (control plane); cPCI and H.110 compatible
- Overlays a packet-based switching architecture on top of CPCI to create an Embedded System Area Network (ESAN)
- Integrates Ethernet on backplane
- 95% of world-wide data travels on Ethernet! (up to 2Gb/s full duplex)
- 85% of installed networks are Ethernet
  Ethernet technology continues to be incorporated into more products than ever before
- Standard cPCI connector

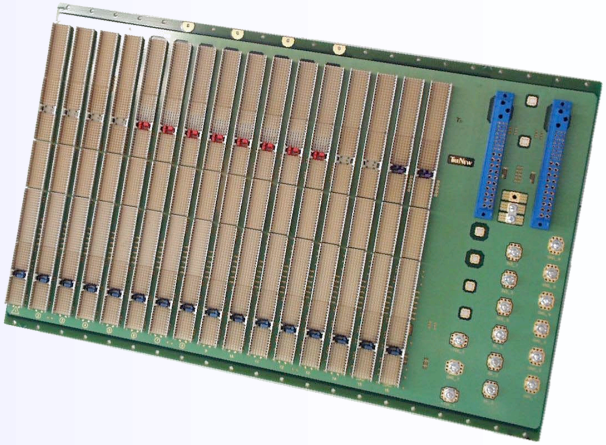

PICMG 2.16

Node and Fabric pin out

- **Connection to Node slots is done via J3/ P3 pins (up to 4 Gbps)**
- **Connection to Fabric Slots is done via P3/P5 pins (up to 40/50 Gbps)**

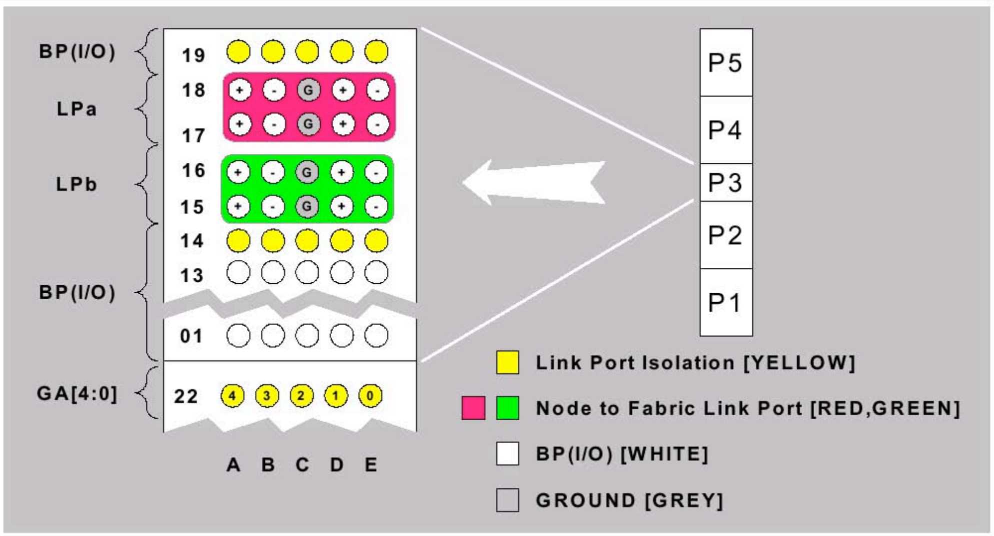

Node Slot Pin Allocation

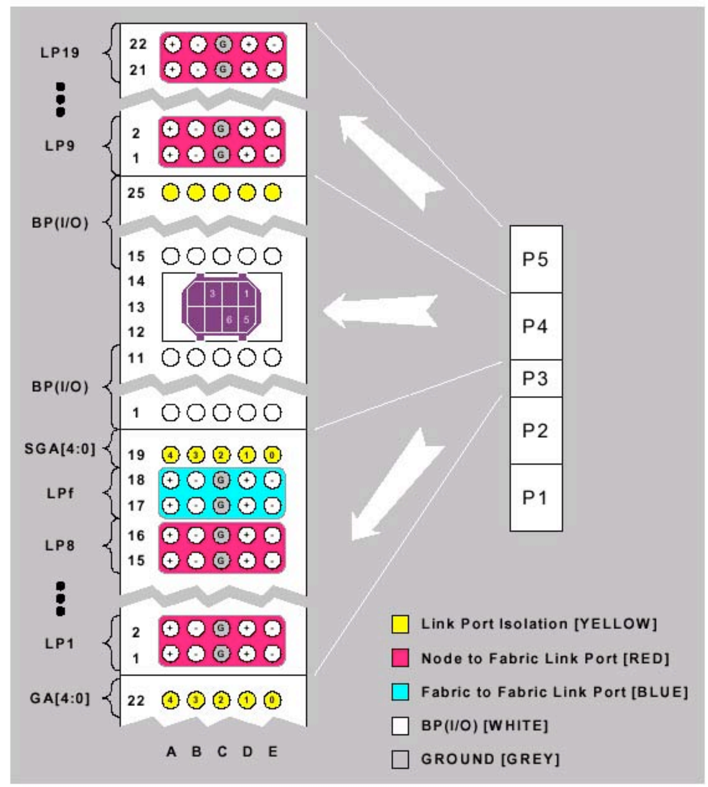

Standard Fabric Slot Pin Allocation

VITA 31.1 Gigabit Ethernet on VME64x

bus tronic
An Elma Company

About VITA 31.1

→**Overview:** Defines a pin out for implementing a 10/100/1000BASE-T Ethernet switched network across the P0/J0 VME64x backplane connectors.

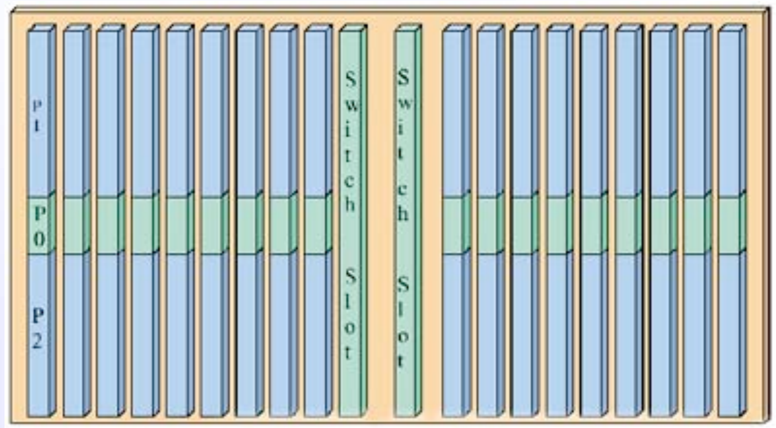

Possible VITA31.1 configuration

→ **Topologies:** Single star fabric switch or dual star fabric switch architecture.

→ **Fabric Slot:** Requires 1 or 2 dedicated fabric-switch slots. These utilize 5 2mm HM connectors and pin assignments are identical to the PICMG 2.16 Packet Switched Bus. Fabric switch cards can be used in either the VME or CompactPCI Environment. node boards (specifies 2mm HM numbering from bottom up as in PICMG)

→ **Node Slot:** Fabric A – B19 P0 rows 2,3; Fabric B- B19 P0 rows 4,5; GNDs P0 Row 1,6

Fast & Giga Ethernet Switch

PICMG2.16 & VITA31.1

5100a

Provides 24 Fast Ethernet, 2 Gigabit Ethernet channels and an high speed bus for network expansion options

Compliant with PICMG 2.16 or VITA 31.1 system

High speed non-blocking layer 2 switch with :

- *Store-and-forward*
- *4000 MAC adresses*
- *Static or Automatic MAC address management*
- *Broadcast filtering*

Auto-negotiation and auto-MDI/ MDIX crossover for true Plug-n-Play

Prevents packet loss with back pressure and IEEE 802.3x flow control

Description

ComEth 5100a is a range of highly integrated layer 2 Ethernet switches that provides 24 10/100TX Ethernet ports and two Gigabits ports. These products are built on a single height 6U board.

Switch Fabric Use in Slot SBCs, 2002

2002 Total: 150.2 Thousand Units

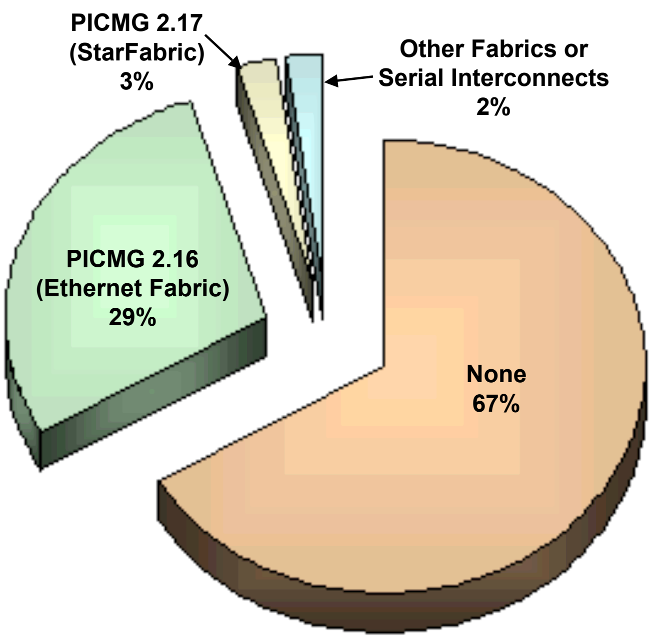

- VDC expects that over 50% of CompactPCI Slot SBCs will be 2.16 compliant in 2004.
- StarFabric has achieved acceptance in several markets, including Military / Aerospace, with several design wins. Product rollouts continue.
- PCI-Express Advanced Switching is expected to have a significant impact when silicon is readily available.
- InfiniBand is alive and well, particularly in "supercomputing"; Embedded InfiniBand initiative under way within IBTA and VITA. Silicon readily available.
- RapidIO product rollouts continue; expect increasing native support for RapidIO in PowerPC CPUs. Xilinx shipping FPGAs containing PowerPC processors and RapidIO interfaces.

Modern Backplane Topologies

Full Mesh

- In a Full Mesh each node has a direct connection to every other node
- Bandwidth aggregates
  - In this example each of 8 nodes has 7 separate connections
  - Assuming 10Gbps links this would yield 8 X 7 X 10Gbps or 560Gbps of aggregate system bandwidth
- Failure overhead is reduced to only n+1

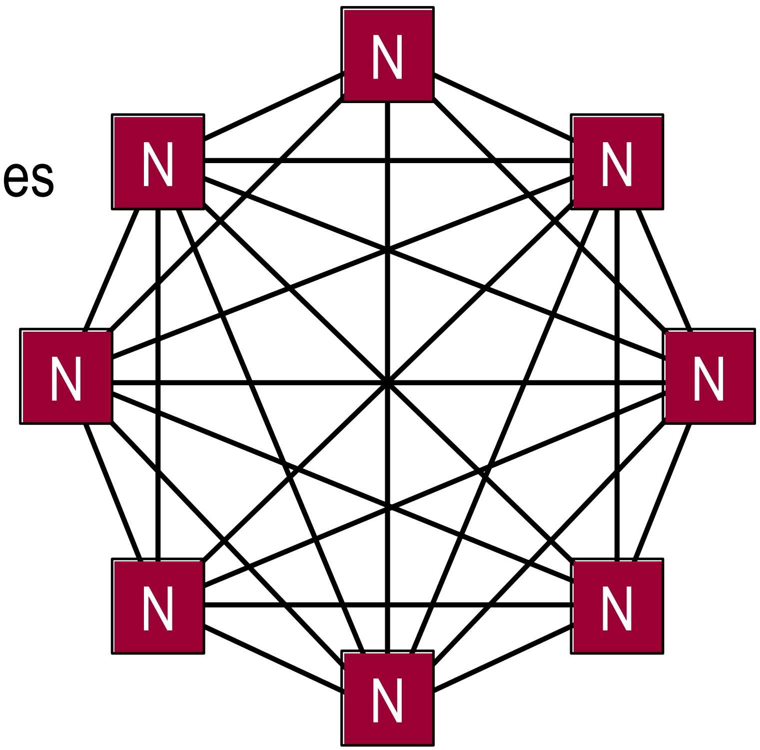

PICM 2.17 StarFabric

About PICMG 2.17

- → Defines backplane, Node Boards and Switch Boards
  - → 3 types of Boards
    - → Basic
    - → Multi-Segment
    - → Fabric Native

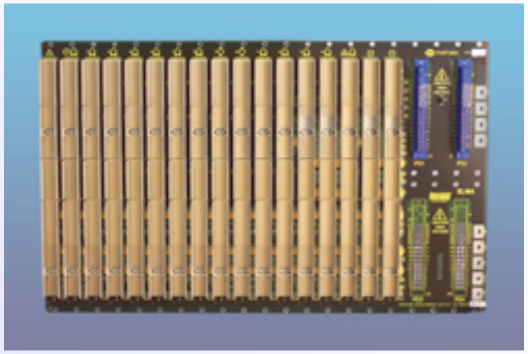

- → Compatible with existing PICMG specs.
- → Migration path to PCI Express Advanced Switching
- → 2.5Gbs per Slot (with four 622Mbytes/s links); roadmap up to 10Gbs
- → High Availability, Multiple classes of traffic (asynchronous and isochronous).
- → Redundant, point to point, high-speed serial connection
- → Virtual Backplane possible (shelf address)
- → CompactPCI connector with the addition of a type A 2mm HM module.

VXS – VITA 41

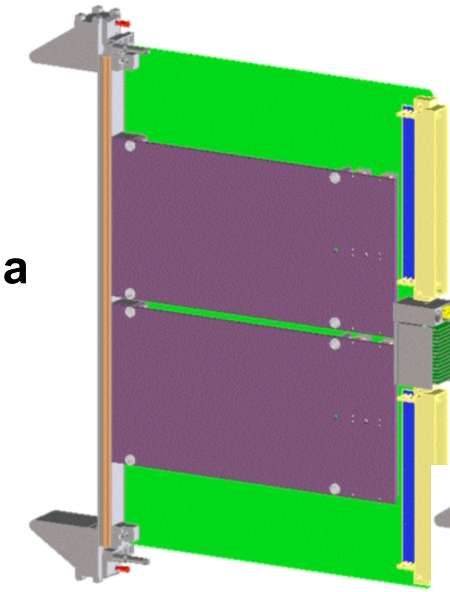

- **Very high speed switched seria interconnect for VME**
- **VXS is a second wave VME Renaissance technology**
- **Two new versions of the spec were ratified during year 2003**
  - **Accommodate conduction cooling**
  - **Accommodate PMC seated in front of new P0 connector**
- **Rear Transition Module (RTM) spec in draft**

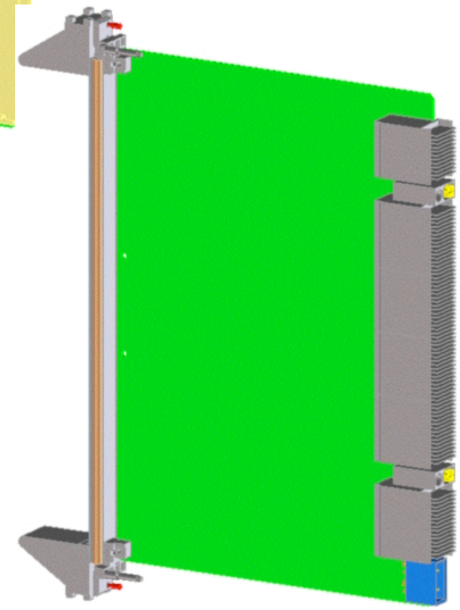

VMEbus Technology™

VXS Backward Compatibility

VME Renaissance

PICMG 1.3 (Industrial PCI Express)

bus tronic
An Elma Company

About PICMG 1.3 (IPCI-E)

→PCI-Express over PICMG 1.X type of backplane

→ Backplane can operate at 1X, 4X, 8X, or 16X

→ Both full and half size SHB (Half-size SHB form factor based on half-size PCI-E board)

→ The full-size SHB and backplane has 16X PCI-Express connectors (A and C) and one 8X PCI-Express connector (B).

→Optional ATX, EPS, or BTX Power Supply Support

VITA 48 Enhanced Ruggedized Design Implementation (ERDI)

About VITA 48

→ Enclosed front removable slot card modules:

- → High power
- → EMI shielding
- → forced air or liquid flow-thru cooling
- → cold plate with ATR clamp-guides
- → rugged modular slot-bay construction

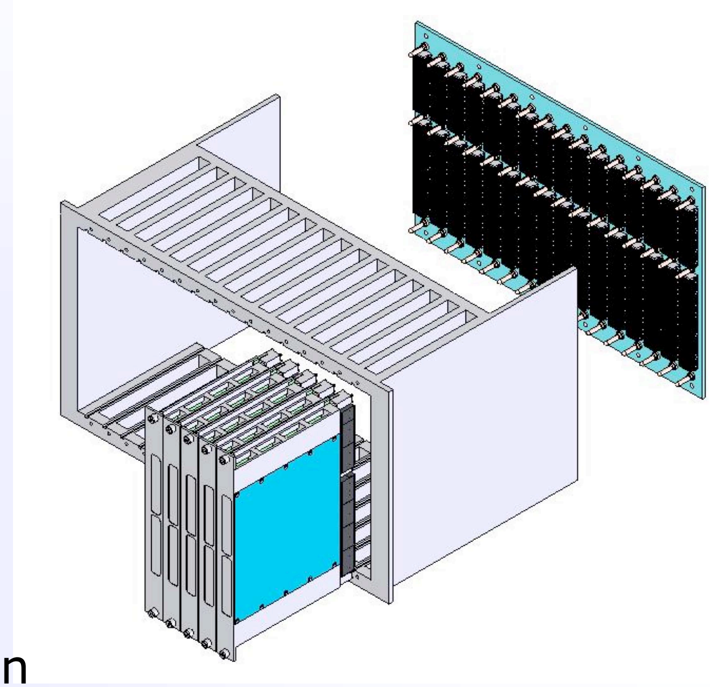

→ Nominal 6U-160mm form factor

→ Slot pitch - 1.00”

CC121 High-Density CompactPCI Crate

General Specifications

Input Voltage Range
85 - 264 VAC
Universal

Voltage Frequency Range
47 – 440 Hz
Auto-frequency ranging

Power Factor
0.99 typical,
meets EN61000-3-2

Slots (6U)	Available Power	Size, w x d x h	Rack Mount Space Needed	Weight
21	1260 W	483 x 440 x 399 mm (411 including the feet)	9U (399 mm)	19.9 kg

	+3.3 VDC	+5 VDC	+12 VDC	-12 VDC
Maximum current	120 A	120 A	50 A	25 A

The CC121 crate accepts modules compliant with CompactPCI, PICMG 2.0 R3.0. The crate accepts PXI modules, but does not provide PXI-specific features.

The CC121 does not accept rear panel transition boards.

Warranty
3 years

Environmental

Operating Temperature
0° to 45°C

Operating Relative Humidity
5 to 95% non-condensing

Shock & Vibration*
30 G half-sine pulse,
5 – 500 Hz random

EMC Immunity
Complies with EN61326-1 industrial Environment

EMC Emissions
Complies with EN61326-1, EN55011Class B

Mechanical Packaging
Complies with IEEE1101.10

Safety
Complies with EN61010-1

safety class 1

CE Certification and Compliance

* As defined by MIL-PRF-28800F Class 3

CC121 equipped with 20 1-4 GS/s quad-channel Digitizers (DC271) and one PCI interface for easy connection to a benchtop PC

NALLATECH

BenADIC™

CompactPCI 20 Channel ADC Motherboard

Simultaneous 20-Channel, 14-bit ADC, operating at 105MSPS with 250Mhz Bandwidth

Overview

The BenADIC provides high performance Analogue-to-Digital Data Conversion on a cPCI platform. High performance on the BenADIC is achieved by an array of 20 ADCs tightly integrated into an FPGA network.

ADCs

Each 14-bit ADC operates concurrently and can be accessed via the front panel. Each ADC on the BenADIC operates at 15-105MSPS and supports either a single-ended or differential analogue input.

BenADIC Architecture

Seven FPGAs on the BenADIC provide maximum flexibility when interfacing with either the ADCs, PCI Interface, or backplane connectors. To further enhance the FPGAs an extensive communications infrastructure has been designed between each FPGA, providing distributed and point to point parallel architectures.

BenADIC

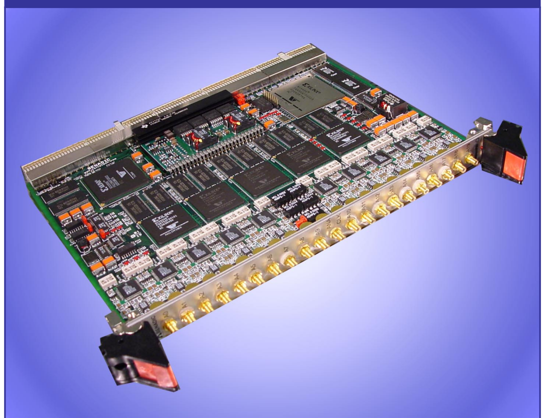

VIRTEX VIRTEX-E

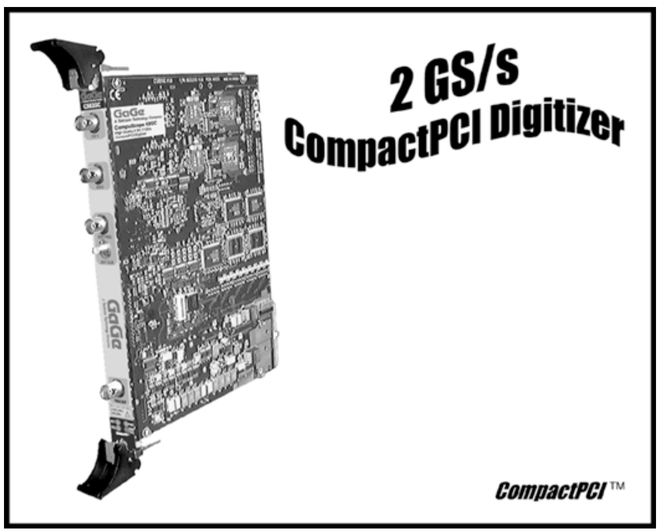

PXD Series

PXI Digitizers

Leading Features

- *Complete family from 150 MHz to 1 GHz bandwidths*
- *Up to 2 GS/s sample rate*
- *Up to 8 million points of acquisition memory*
- *Up to 50 GS/s Random Interleaved Sampling (RIS) for repetitive applications*

The complete line of eight LeCroy PXD Digitizers from 150 MHz to 1 GHz Bandwidths

CAEN

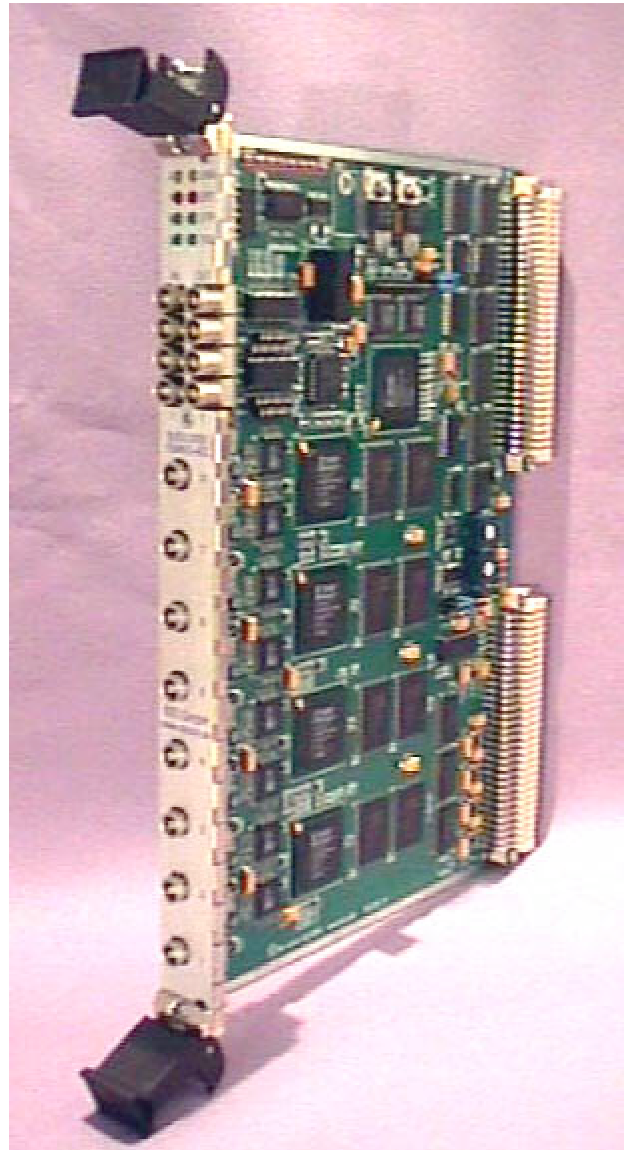

Struck

Integrating High Performance VME, PCI and CPCI processors into CPU clusters for Data Acquisition and Control Systems

L. Vivolo, M. Weymann, F. H. Worm

Creative Electronic System,
70 Route du Pont-Butin,
CH 1213 Petit-Lancy/Geneva,Switzerland
email: `ces@ces.ch`, web: `http://www.ces.ch`

Abstract

Data acquisition and control systems using a large number of embedded VME processors have a long tradition in High Energy Physics. More recently, CPCI is in some cases considered as an alternative to VME. CES has developed a processor board architecture optimized for high-throughput deterministic bus operation, which can used with minimal adaptation on both backplanes. The RIO3 8064 (VME) and RIOC 4065 (CPCI) are to a large extend software compatible which helps to develop software solutions almost simultaneaously in both domains. Both boards couple the CPU bus directly to the backplane bus (VME or CPCI) and to two independent PCI busses. This twin-bus architecture allows to separate data flows in a similar way than VME/VSB architectures did in the past. The MFCC 844x, a PowerPC based PMC module is ideally suited to handle complex I/O protocols or to build multi-CPU clusters coupled by the carrier boards local PCI bus rather than the backplane bus. The CES PVIC allows to interconnect distant PCI segments (e.g a VME based processor cluster and a desktop workstation) using both memory mapped access and DMA mechanisms. With its *backplane driver* CES provides an ideal tool to integrate CPUs interconnected by PCI,VME,CPCI and PVIC into a homogeneous, network-oriented environment taking full advatage of the high bandwidth and low latency features of PCI and bakplane busses.

CPCI PROCESSOR BOARDS

RIOC 4067

PowerPC-Based CompactPCI 2.16 Switched Backplane Real-Time Processor Board

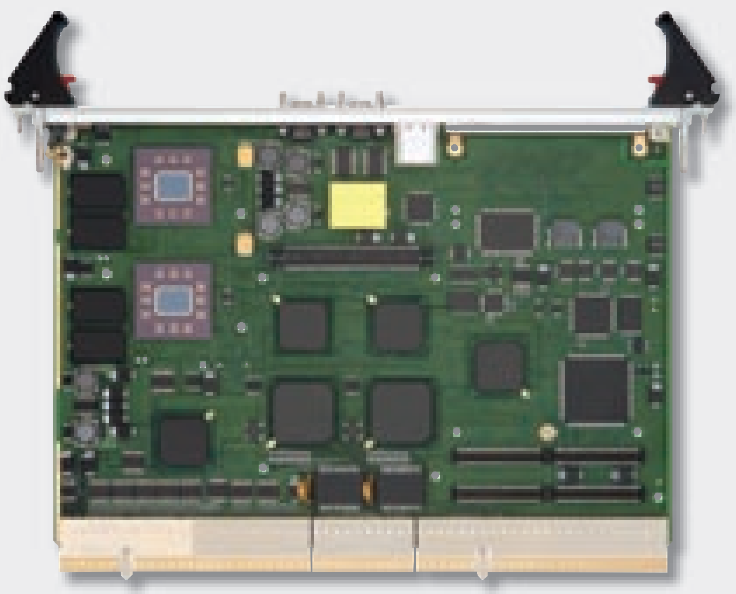

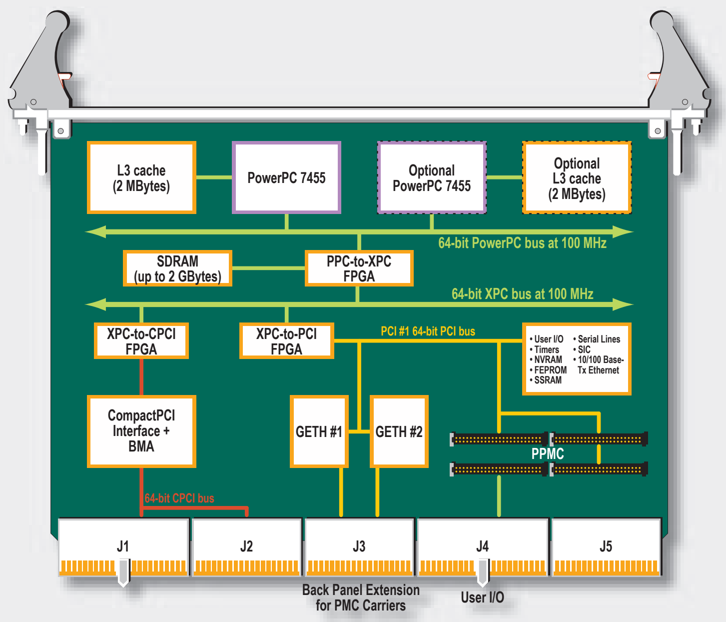

CES - CREATIVE ELECTRONIC SYSTEMS

GENERAL PURPOSE I/Os

GPIO 8409

Reconfigurable I/O PMC and Rocket I/O Interconnect PMC

Reconfigurable !

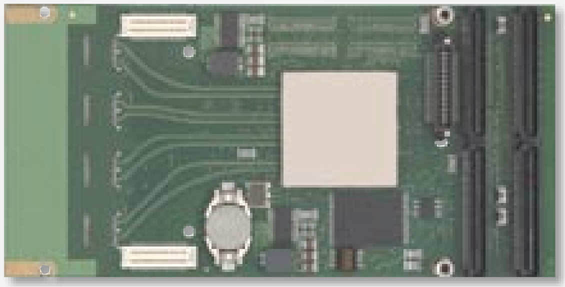

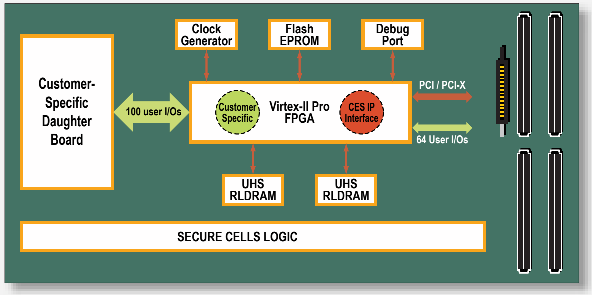

GENERAL DESCRIPTION

The GPIO 8409 has been designed for building I/O subsystems connected to the PCI bus of the VME and CompactPCI CES real-time platforms.

The front-end section incorporates an ultra-high-speed large dimension programmable logic device (FPGA Virtex-II Pro XC2VP20 / XC2VP30), which is connected on one side to the PCI bus and on the other side either to the front panel via an electrical adaptor or to the PMC I/O connector (PN4) on the rear.

Due to the removable front-end adaptor and to the reprogrammable FPGA, the GPIO can be used to execute different I/O functions at different times as user-requirements change.

An extended specification and a conduction cooled version will be available soon.

GPIO 8409CA - ROCKET I/O INTERCONNECT

The GPIO 8409CA supports the RocketIO connections with front panel small form package connectors, which can house a variety of optical transceivers. It also maintains 100 independent user-accessible I/O buses.

- 64 MBytes of ultra-high-speed RLDRAM
- 32 MBytes Flash EPROM
- Onboard Virtex-II Pro FPGA
- Houses four pluggable SFP transceivers
- Supports single-mode and multi-mode fibers at up to 3.125 Gbit/s
- 32 / 64-bit PCI or PCI-X interface
- Multiple thermal sensors
- Secure cells package (option)
- Sockets for user-specific mini mezzanine

Data Acquisition for **DCBA**

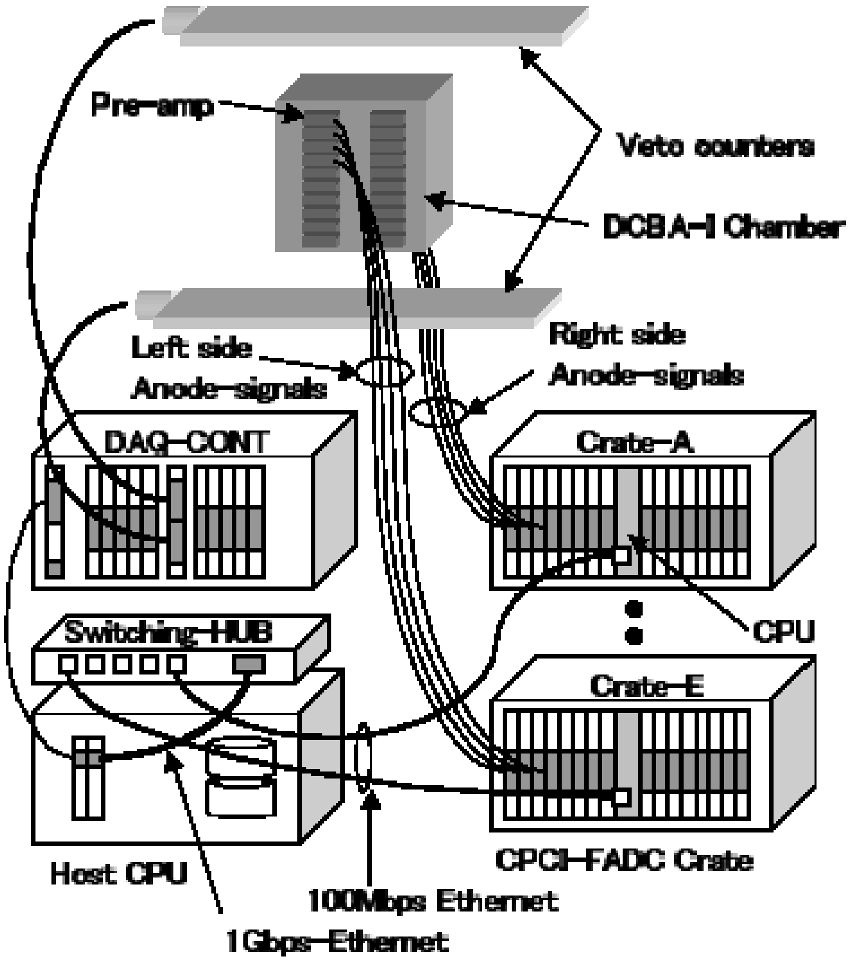

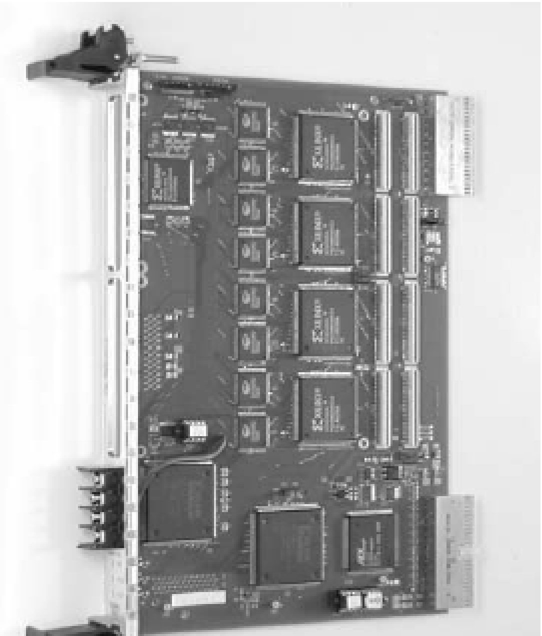

Digital

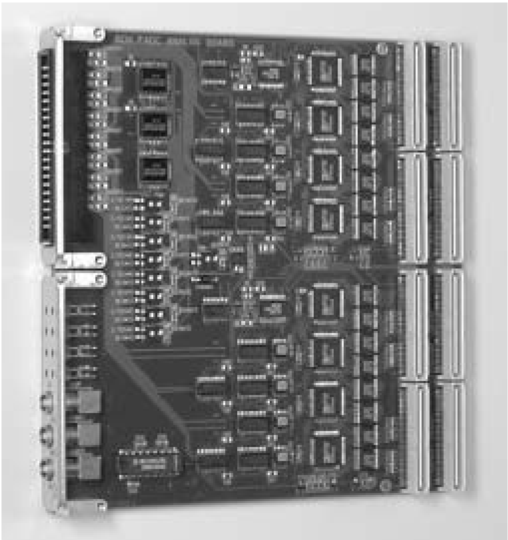

Analog

Recommendations (Opinions):

- IU: Continue prototype of cPCI fADC with “parallel” energy sum for calorimeters
- CNU: Continue prototype study of “serial” energy sum
- JLab: Continue “universal” fADC design, test in existing halls
- Don’t overconstrain GlueX design space yet
- Keep an eye on emerging “form factors”
- CODA group should consider ramifications of “all ethernet” solution
- GlueX will have electronics/DAQ collaborators from outside lab

“push” vs “pull”

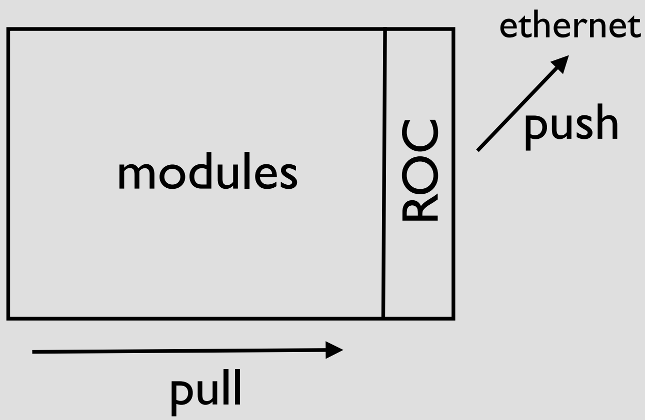

Current CODA model

sparse data
event interrupt
buffer in ROC

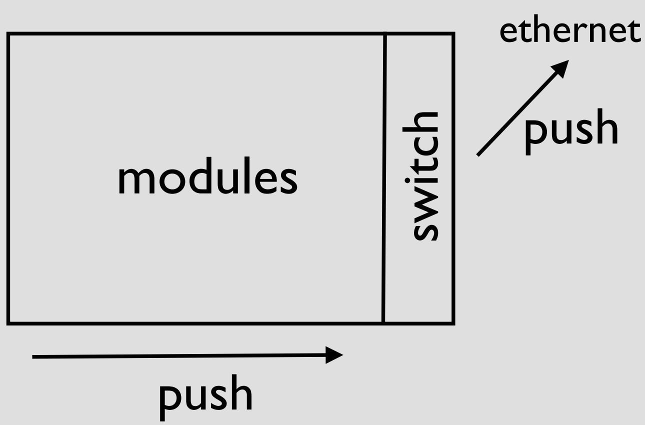

GlueX CODA model?

large data blocks
no event interrupt
buffer in modules