IU GlueX electronics R&D status

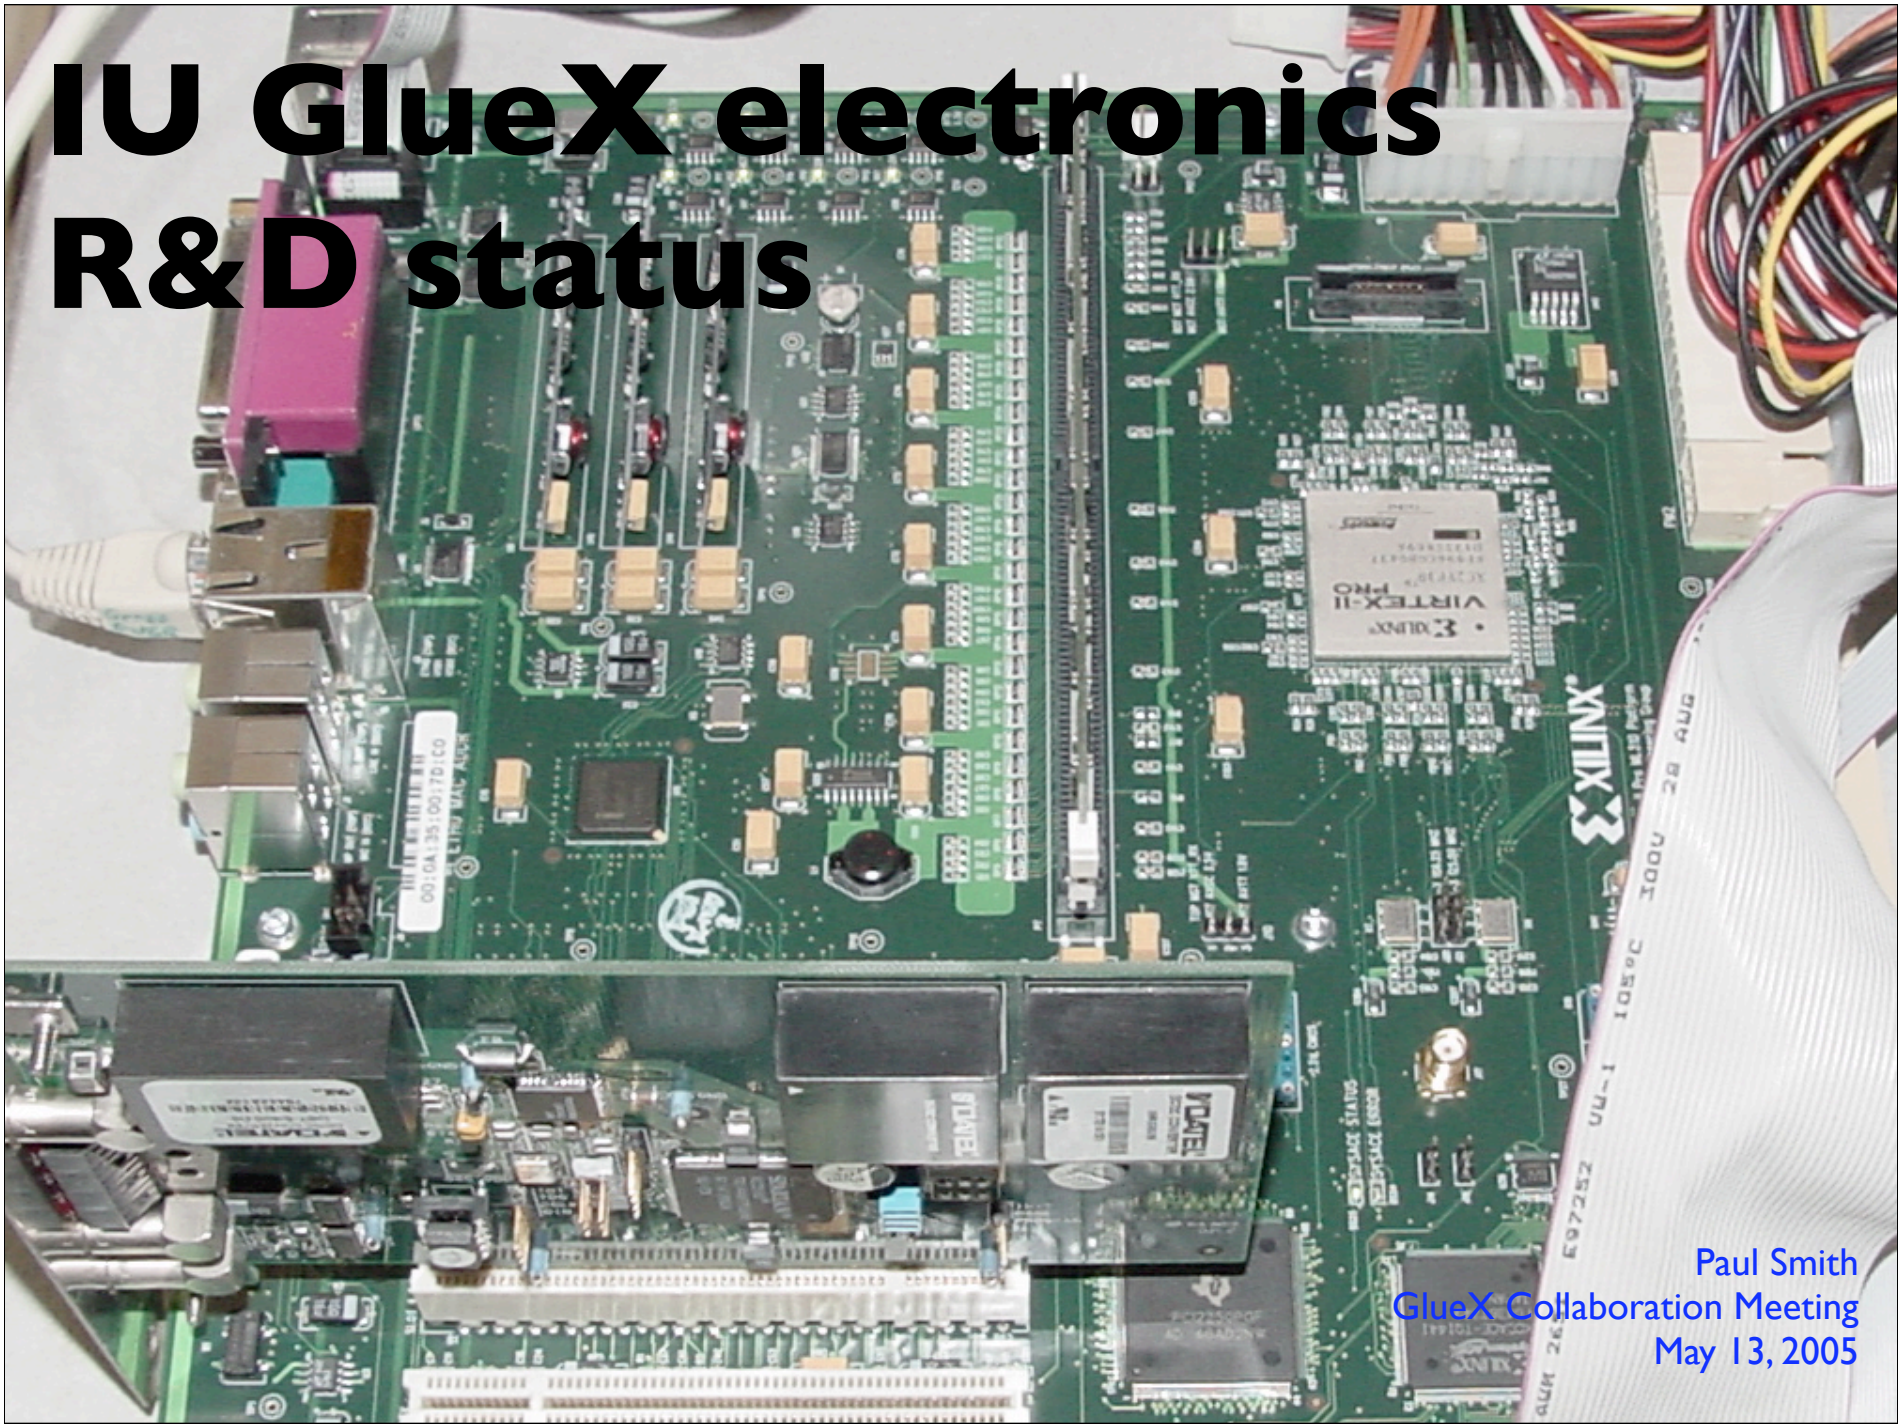

Paul Smith
GlueX Collaboration Meeting
May 13, 2005

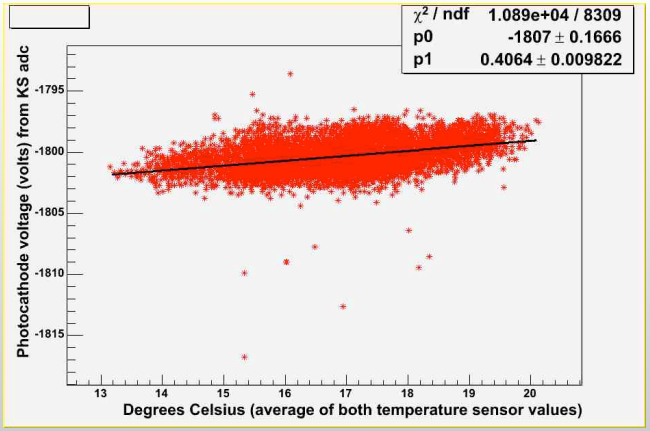

FCal Cockcroft-Walton PMT base testing

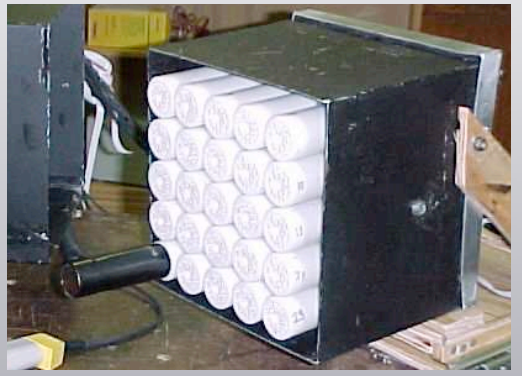

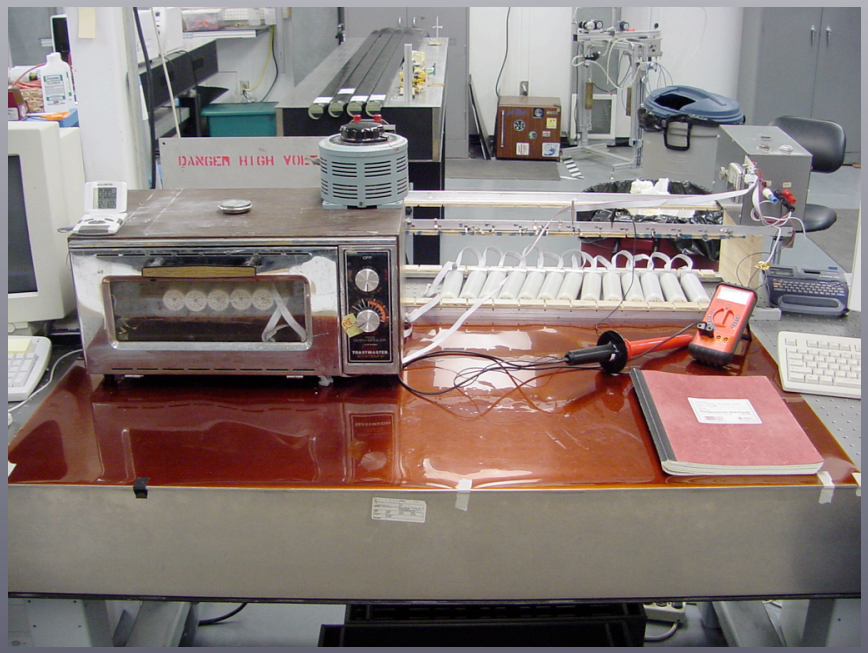

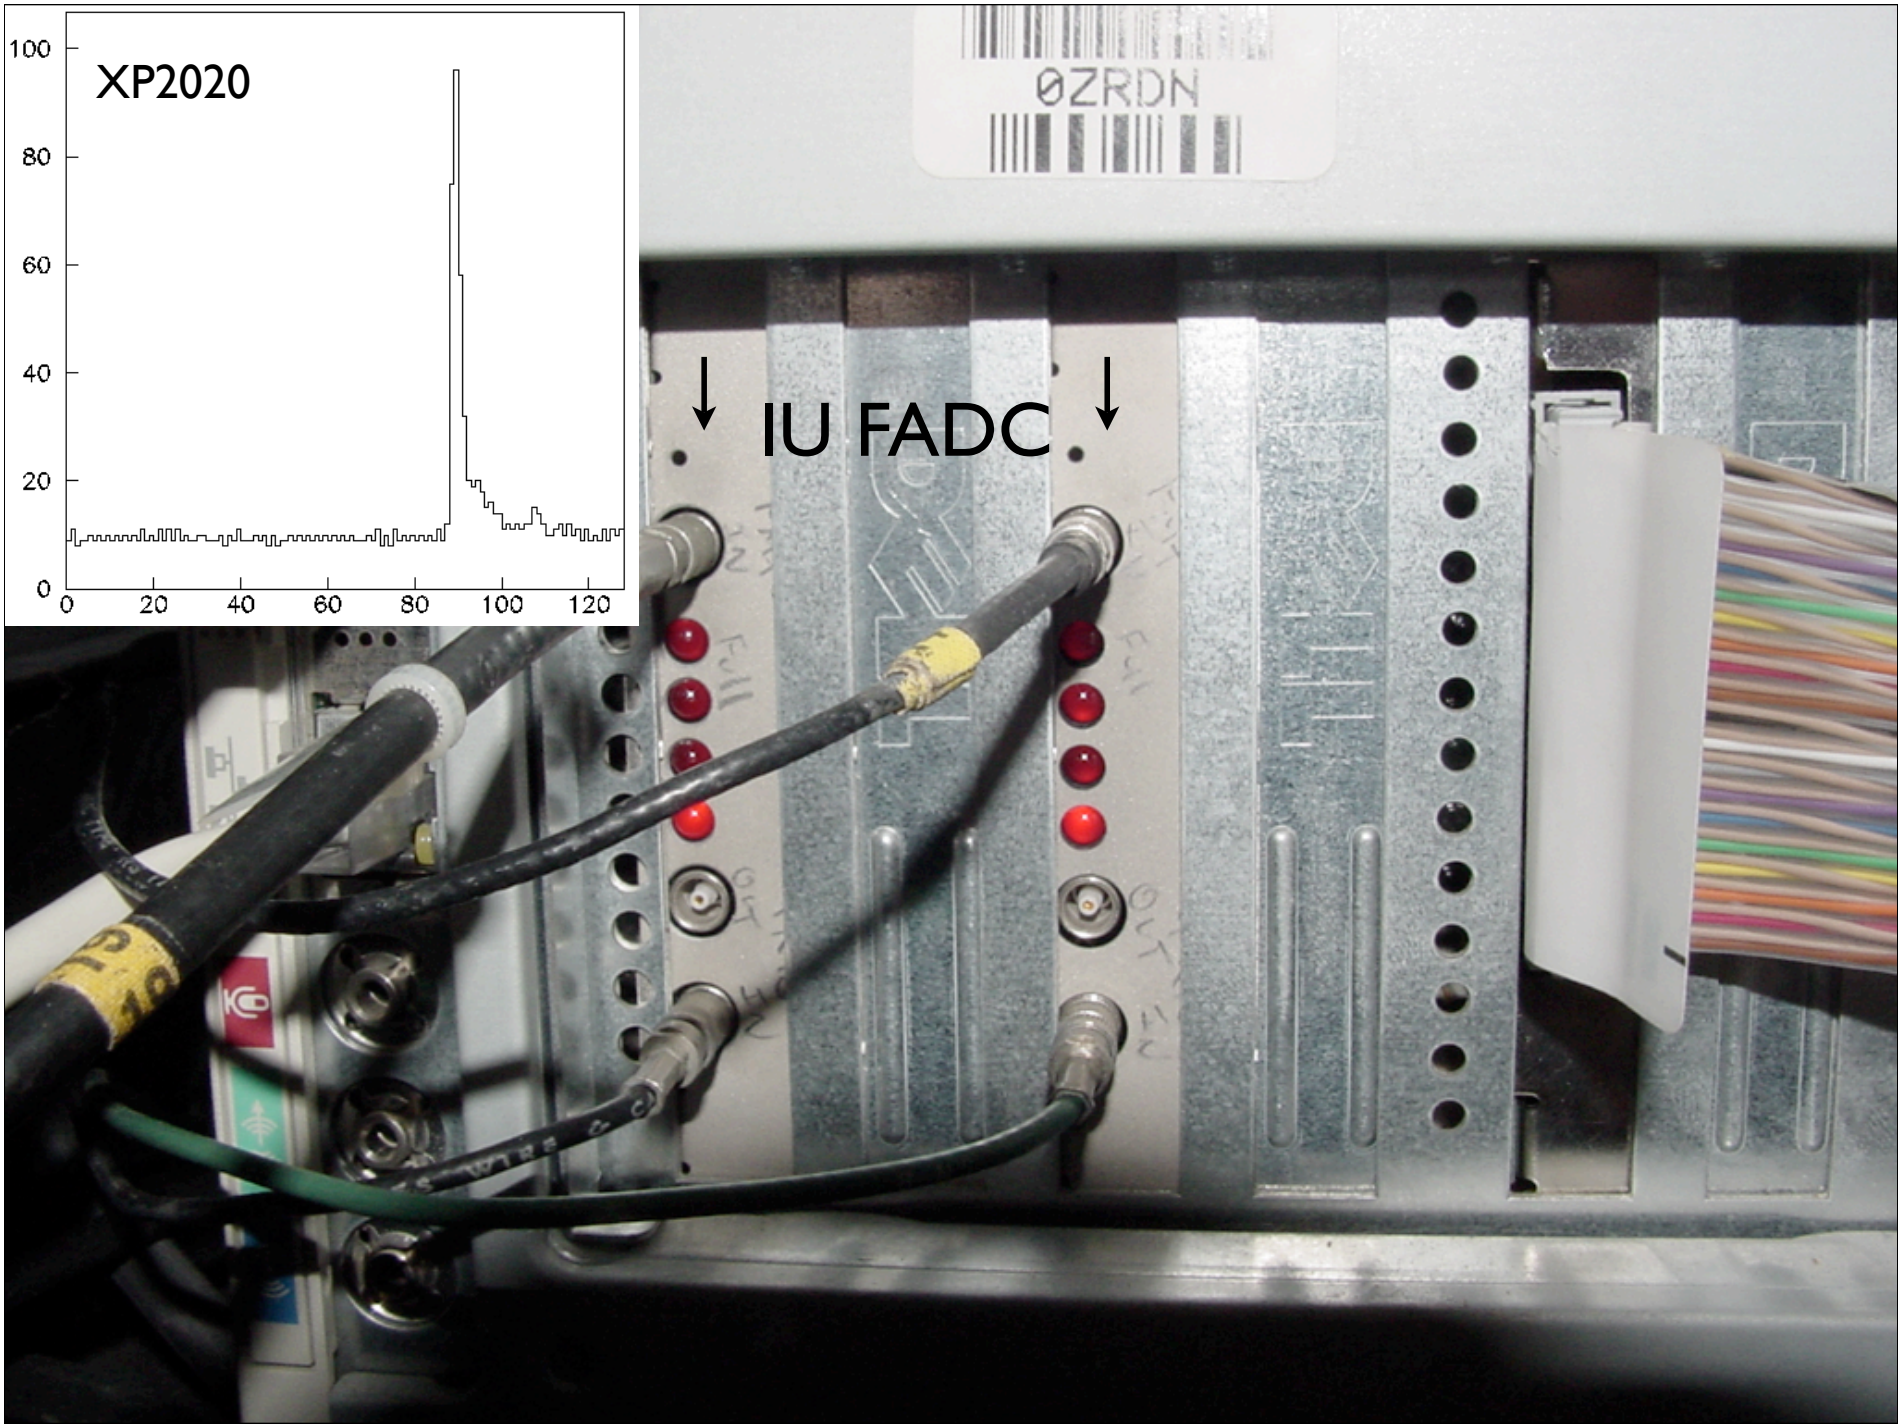
XP2020
100
80
60
40
20
0
0
20
40
60
80
100
120
IU FADC
0ZRDN

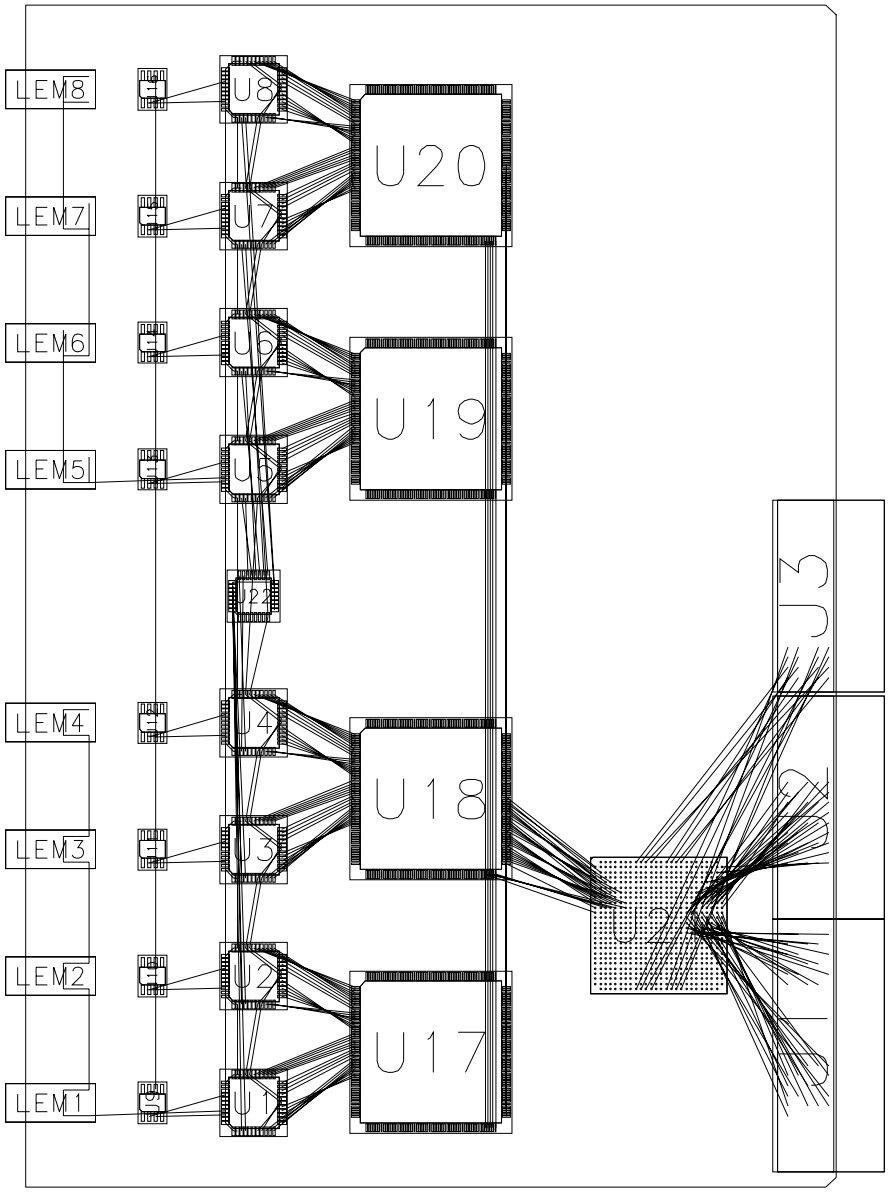

AD9480 $18 (1000)

XC3S50 $12 (100)

XC4VFX12 $120 (100)

8 channel cPCI Calorimeter fADC prototype

Per channel budget:	
Lemo	$5
Amp	$5
Digitizer	$18
Ring Buffer, FIFO	$6
Clocking	$5
DACs	$5
Board, connectors, power	$320/8=$40
PPC gate array	$120/8=$15
SDRAM	$64/8=$8
boot Flash, controller	$96/8=$12
ethernet PHY	$48/8=$6
total:	$125

Why XC4VFX12?

12K logic cells

81 kB dual port RAM

PPC 405 with APU

32 multipliers

2 ethernet MACs

linux

algorithm partitioning

reprogrammable

source synchronous

i/o for energy sum

XILINX®

Virtex-4 Family Overview

DS112 (v1.2) December 8, 2004 **Advance Product Specification**

General Description

The Virtex-4™ Family is the newest generation FPGA from Xilinx. The innovative Advanced Silicon Modular Block or ASMBL™ column-based architecture is unique in the programmable logic industry. Virtex-4 FPGAs contain three families (platforms): LX, FX, and SX. Choice and feature combinations are offered for all complex applications. A wide array of hard-IP core blocks complete the system solution. These cores include the PowerPC™ processors (with a new APU interface), Tri-Mode Ethernet MACs, 622 Mb/s to 11.1 Gb/s serial transceivers, dedicated DSP slices, high-speed clock management circuitry, and source-synchronous interface blocks. The basic Virtex-4 building blocks are an enhancement of those found in the popular Virtex-based product families: Virtex, Virtex-E, Virtex-II, Virtex-II Pro, and Virtex-II Pro X, allowing upward compatibility of existing designs. Virtex-4 devices are produced on a state-of-the-art 90-nm copper process, using 300 mm (12 inch) wafer technology. Combining a wide variety of flexible features, the Virtex-4 family enhances programmable logic design capabilities and is a powerful alternative to ASIC technology.

One or Two PowerPC 405 Processor Cores

- 32-bit Harvard Architecture
- 5-Stage Execution Pipeline
- Integrated 16KB Level 1 Instruction Cache and 16KB Level 1 Data Cache
  - Integrated Level 1 Cache Parity Generation and Checking
- CoreConnect™ Bus Architecture
- Efficient, high-performance on-chip memory (OCM) interface to block RAM
- PLB Synchronization Logic (Enables Non-Integer CPU-to-PLB Clock Ratios)
- Auxiliary Processor Unit (APU) Interface and Integrated APU Controller
  - Optimized FPGA-based Coprocessor connection
    - Automatic decode of PowerPC floating-point instructions
  - Allows custom instructions (Decode for up to eight instructions)
  - Extremely efficient microcontroller-style interfacing

SelectIO Technology

- Up to 960 user I/Os
- Wide selections of I/O standards from 1.5V to 3.3V
- Extremely high-performance
  - 600 Mb/s HSTL & SSTL (on all single-ended I/O)
  - 1 Gb/s LVDS (on all differential I/O pairs)
- True differential termination
- Selected low-capacitance I/Os for improved signal integrity
- Same edge capture at input and output I/Os
- Memory interface support for DDR and DDR-2 SDRAM, QDR-II, RLDRAM-II, and FCRAM-II

ChipSync Technology

- Integrated with SelectIO technology to simplify source-synchronous interfaces
- Per-bit deskew capability built in all I/O blocks (variable input delay line)
- Dedicated I/O and regional clocking resources (pin and trees)
- Built in data serializer/deserializer logic in all I/O and clock dividers
- Memory/Networking/Telecommunication interfaces up to 1 Gb/s+

Two or Four Tri-Mode (10/100/1000 Mb/s) Ethernet Media Access Control (MAC) Cores

- IEEE 802.3-2000 Compliant
- MII/GMII Interface or SGMII (when used with RocketIO Transceivers)
- Can Operate Independent of PowerPC processor
- Half or Full Duplex
- Supports Jumbo Frames
- 1000 Base-X PCS/PMA: When used with RocketIO MGT can provide complete 1000 Base-X implementation on-chip

```
BusyBox v0.60.5 (2005.02.28-07:38+0000) Built-in shell (msh)
Enter 'help' for a list of built-in commands.

# ifconfig
eth0      Link encap:Ethernet  HWaddr 00:11:0C:00:16:7D
          inet addr:192.168.100.240  Bcast:192.168.100.255  Mask:255.255.255.0
          UP BROADCAST NOTRAILERS RUNNING  MTU:1500  Metric:1
          RX packets:9 errors:0 dropped:0 overruns:0 frame:0
          TX packets:9 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 iB)  TX bytes:0 (0.0 iB)
          Interrupt:30 Base address:0x300

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:0 (0.0 iB)  TX bytes:0 (0.0 iB)

# ls
bin   dev   etc   home  lib   mnt   proc  sbin  tmp   usr   var
# ls home
httpd
# ls etc
config      inetd.conf  inittab     rc.reboot   services
default     init.d      rc.d        rc.sysinit
# free
              total         used         free       shared      buffers
  Mem:         30112         2708        27404            0          720
# free -k
BusyBox v0.60.5 (2005.02.28-07:38+0000) multi-call binary

# cd /
# ls
bin   dev   etc   home  lib   mnt   proc  sbin  tmp   usr   var
# ls bin
busybox    dhcpcd     hostname   mkdir      ps         tinylogin
cat        echo       inetd      mount      pwd        touch
chmod      erase      kill       msh        rm         true
cp         false      ln         mv         sh         uname
date       flatfsd    login      netflash   telnetd    version
dd         ftpd       ls         ping       thttpd     vi
# ls sbin
getty     halt      ifconfig  init      insmod    reboot    rmmod     route
#
```

Welcome to SUZAKU

This is a placeholder page in the SUZAKU, running uClinux release of the thttpd Web server.

Atmark Techno, Inc., Sapporo, Japan
April 18, 2004

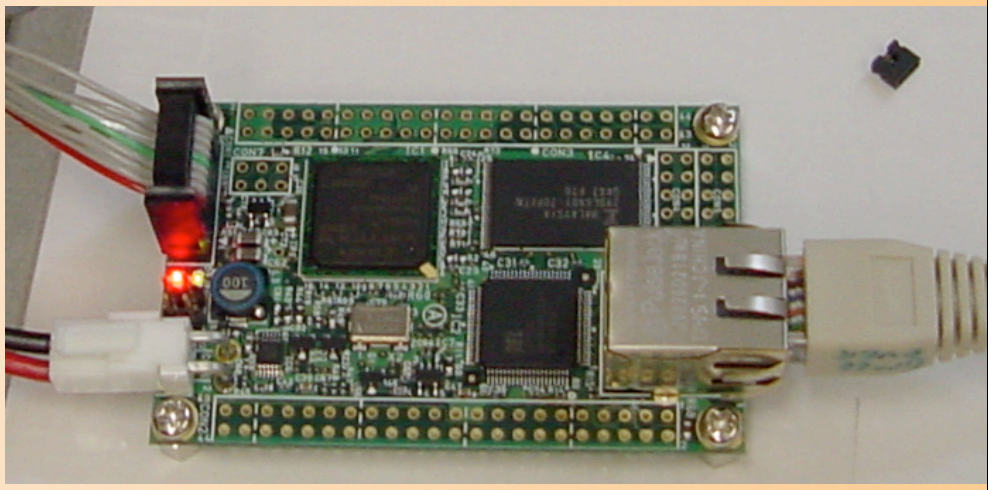

Table 4-1 SUZAKU-V Specifications

FPGA		Xilinx Virtex-□Pro XC2VP4 FG256
Hard Core Processor		PowerPC405
Crystal Oscillator		3.6864MHz (frequency multiplied by FPGA's internal DCM)
Memory	BRAM	16Kbyte
	SDRAM	32Mbyte
	FLASH Memory	8Mbyte
Configuration		Stored on FLASH memory, Controller TE7720
JTAG		2 ports (FPGA and TE7720)
Ethernet		10Base-T／100Base-Tx
Serial		UART 115.2kbps
Timer		2-ch (1-ch is used for OS)
Free I/O Pin		70-pin
Reset Function		Software reset
Power Supply		Voltage: 3.3V±3% Consumption current: 460mA typ (during the operation of a processor)
Board Dimensions		72×47mm

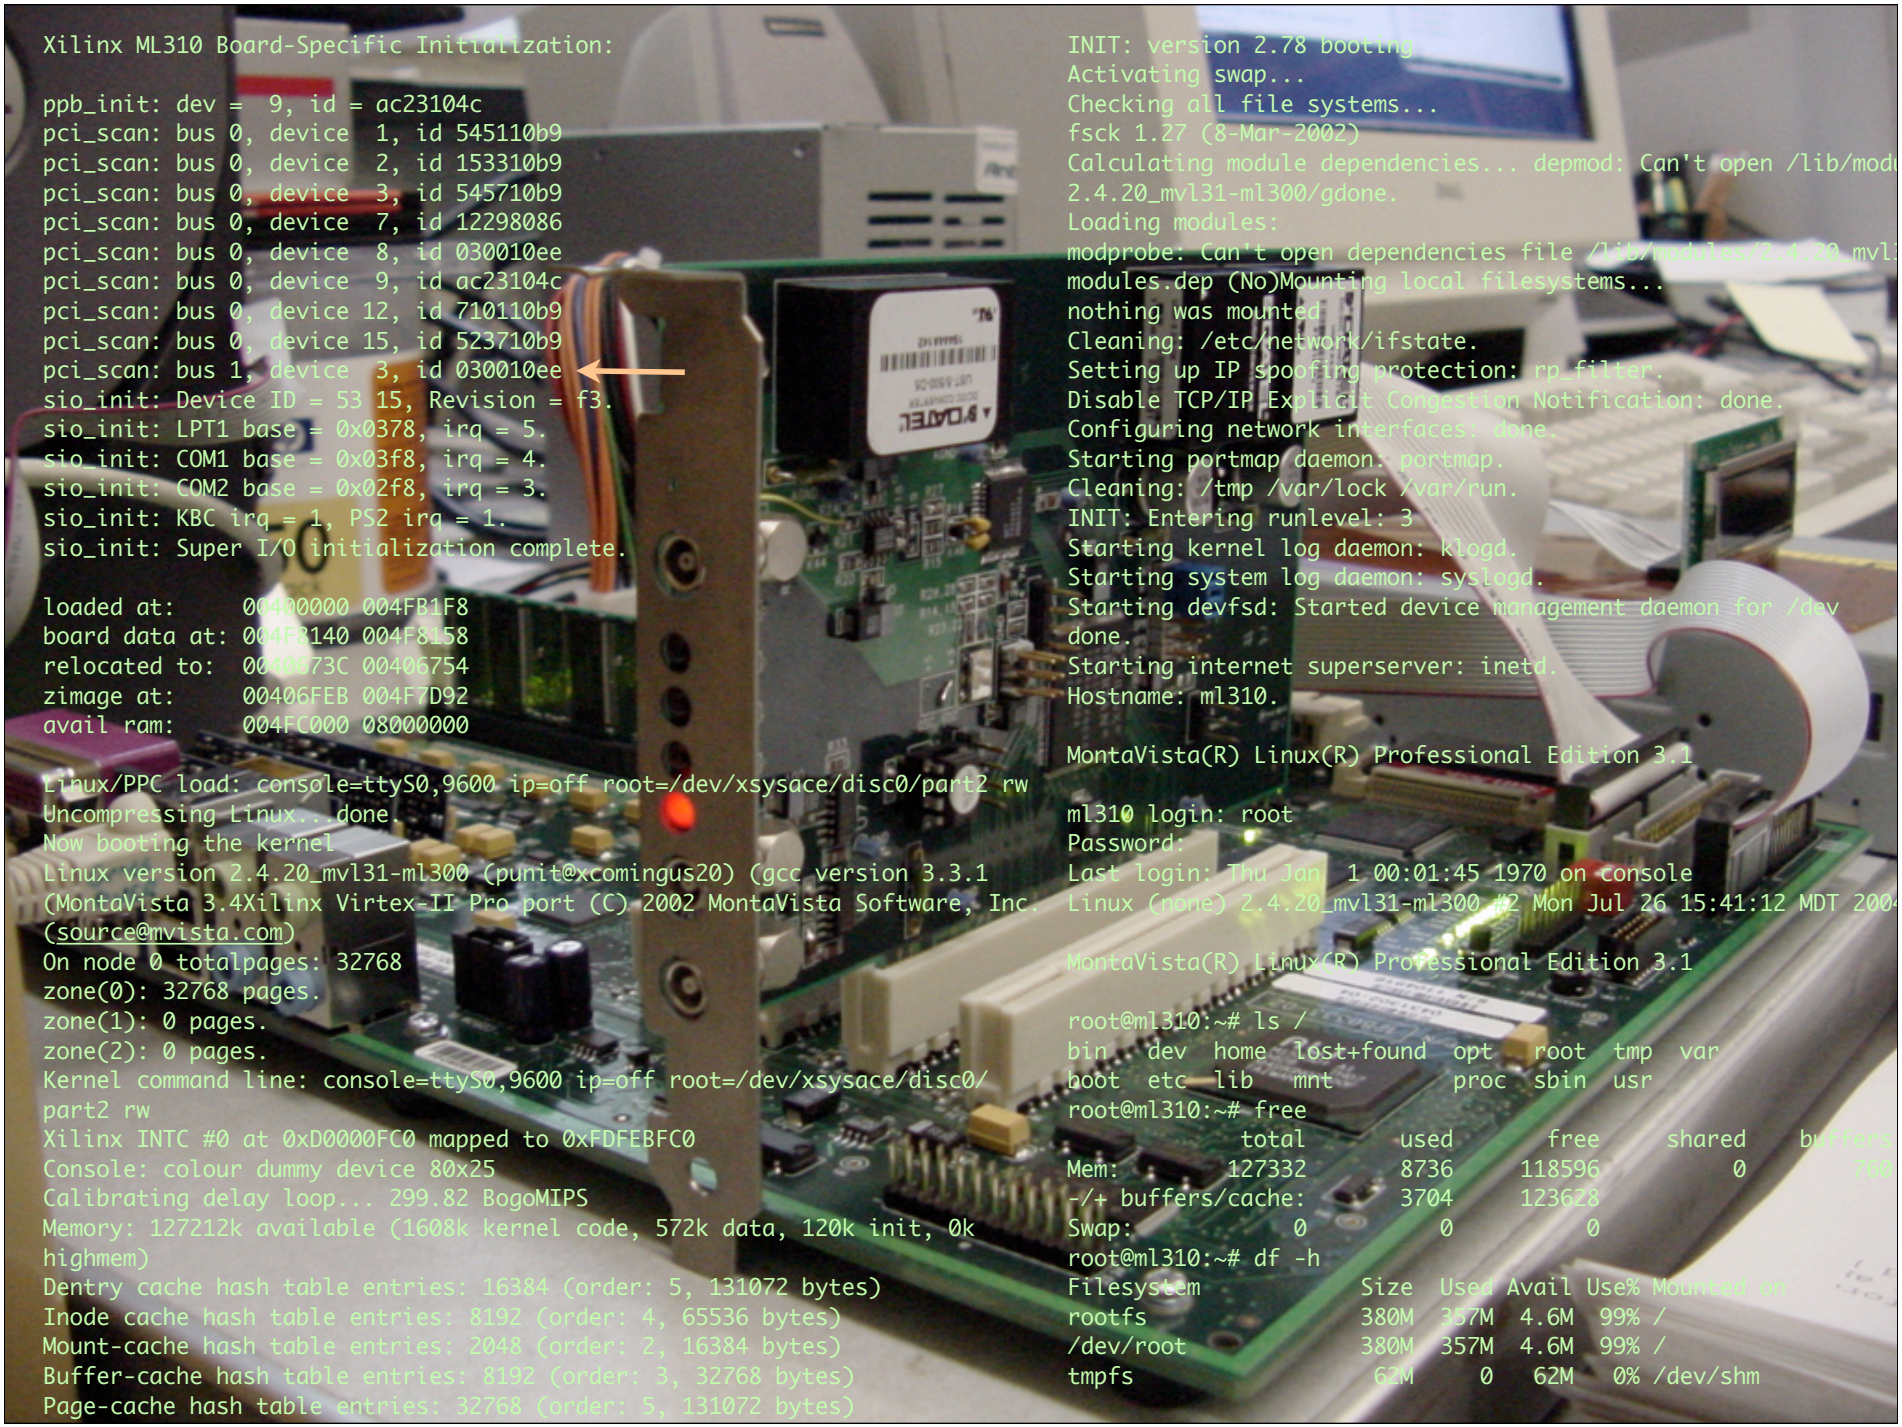
Xilinx ML310 Board-Specific Initialization:
ppb_init: dev = 9, id = ac23104c
pci_scan: bus 0, device 1, id 545110b9
pci_scan: bus 0, device 2, id 153310b9
pci_scan: bus 0, device 3, id 545710b9
pci_scan: bus 0, device 7, id 12298086
pci_scan: bus 0, device 8, id 030010ee
pci_scan: bus 0, device 9, id ac23104c
pci_scan: bus 0, device 12, id 710110b9
pci_scan: bus 0, device 15, id 523710b9
pci_scan: bus 1, device 3, id 030010ee
sio_init: Device ID = 53 15, Revision = f3.
sio_init: LPT1 base = 0x0378, irq = 5.
sio_init: COM1 base = 0x03f8, irq = 4.
sio_init: COM2 base = 0x02f8, irq = 3.
sio_init: KBC irq = 1, PS2 irq = 1.
sio_init: Super I/O initialization complete.
loaded at: 00400000 004FB1F8
board data at: 004F8140 004F8158
relocated to: 0040673C 00406754
zimage at: 00406FEB 004F7D92
avail ram: 004FC000 08000000
Linux/PPC load: console=ttyS0,9600 ip=off root=/dev/xsysace/disc0/part2 rw
Uncompressing Linux... done
Now booting the kernel
Linux version 2.4.20_mvl31-ml300 (punit@xcomingus20) (gcc version 3.3.1
(MontaVista 3.4Xilinx Virtex-II Pro port (C) 2002 MontaVista Software, Inc.
(source@mvista.com)
On node 0 totalpages: 32768
zone(0): 32768 pages.
zone(1): 0 pages.
zone(2): 0 pages.
Kernel command line: console=ttyS0,9600 ip=off root=/dev/xsysace/disc0/
part2 rw
Xilinx INTC #0 at 0xD0000FC0 mapped to 0xFDFEBFC0
Console: colour dummy device 80x25
Calibrating delay loop... 299.82 BogoMIPS
Memory: 127212k available (1608k kernel code, 572k data, 120k init, 0k
highmem)
Dentry cache hash table entries: 16384 (order: 5, 131072 bytes)
Inode cache hash table entries: 8192 (order: 4, 65536 bytes)
Mount-cache hash table entries: 2048 (order: 2, 16384 bytes)
Buffer-cache hash table entries: 8192 (order: 3, 32768 bytes)
Page-cache hash table entries: 32768 (order: 5, 131072 bytes)
INIT: version 2.78 booting
Activating swap...
Checking all file systems...
fsck 1.27 (8-Mar-2002)
Calculating module dependencies... depmod: Can't open /lib/modu
2.4.20_mvl31-ml300/gdone.
Loading modules:
modprobe: Can't open dependencies file /lib/modules/2.4.20_mvl
modules.dep (No)Mounting local filesystems...
nothing was mounted
Cleaning: /etc/network/ifstate.
Setting up IP spoofing protection: rp_filter.
Disable TCP/IP Explicit Congestion Notification: done.
Configuring network interfaces: done.
Starting portmap daemon: portmap.
Cleaning: /tmp /var/lock /var/run.
INIT: Entering runlevel: 3
Starting kernel log daemon: klogd.
Starting system log daemon: syslogd.
Starting devfsd: Started device management daemon for /dev
done.
Starting internet superserver: inetd.
Hostname: ml310.
MontaVista(R) Linux(R) Professional Edition 3.1
ml310 login: root
Password:
Last login: Thu Jan 1 00:01:45 1970 on console
Linux (none) 2.4.20_mvl31-ml300 #2 Mon Jul 26 15:41:12 MDT 200
MontaVista(R) Linux(R) Professional Edition 3.1
root@ml310:~# ls /
bin dev home lost+found opt root tmp var
boot etc lib mnt proc sbin usr
root@ml310:~# free
total used free shared buffers
Mem: 127332 8736 118596 0 760
-/+ buffers/cache: 3704 123628
Swap: 0 0 0
root@ml310:~# df -h
Filesystem Size Used Avail Use% Mounted on
rootfs 380M 357M 4.6M 99% /
/dev/root 380M 357M 4.6M 99% /
tmpfs 62M 0 62M 0% /dev/shm

From the review:

- The general concept of local sums at the front-end board level, followed by crate-level sums and subsequent transfer to a central "Global LVl-1" processing area, is sound. A concept and proof-of-principle for crate backplane operation at the required high rate needs to be developed for the CDR. If high-speed serial operation proves challenging, the collaboration should explore possible parallel concepts to lower the bus-speed requirements.

Notes:

Trim pedestals at input buffers

Energy sum doesn't need 21 bit resolution

Probably 8 bits is fine (+ overflow?)

Truncate where?

Need to subtract pedestal - where?

Integrate (sum over time) after final sum

Multiplier on each channel before sum tree

Ability to test & monitor energy sum

Test pulsers at inputs

Calorimeter Energy Sum

On each board:

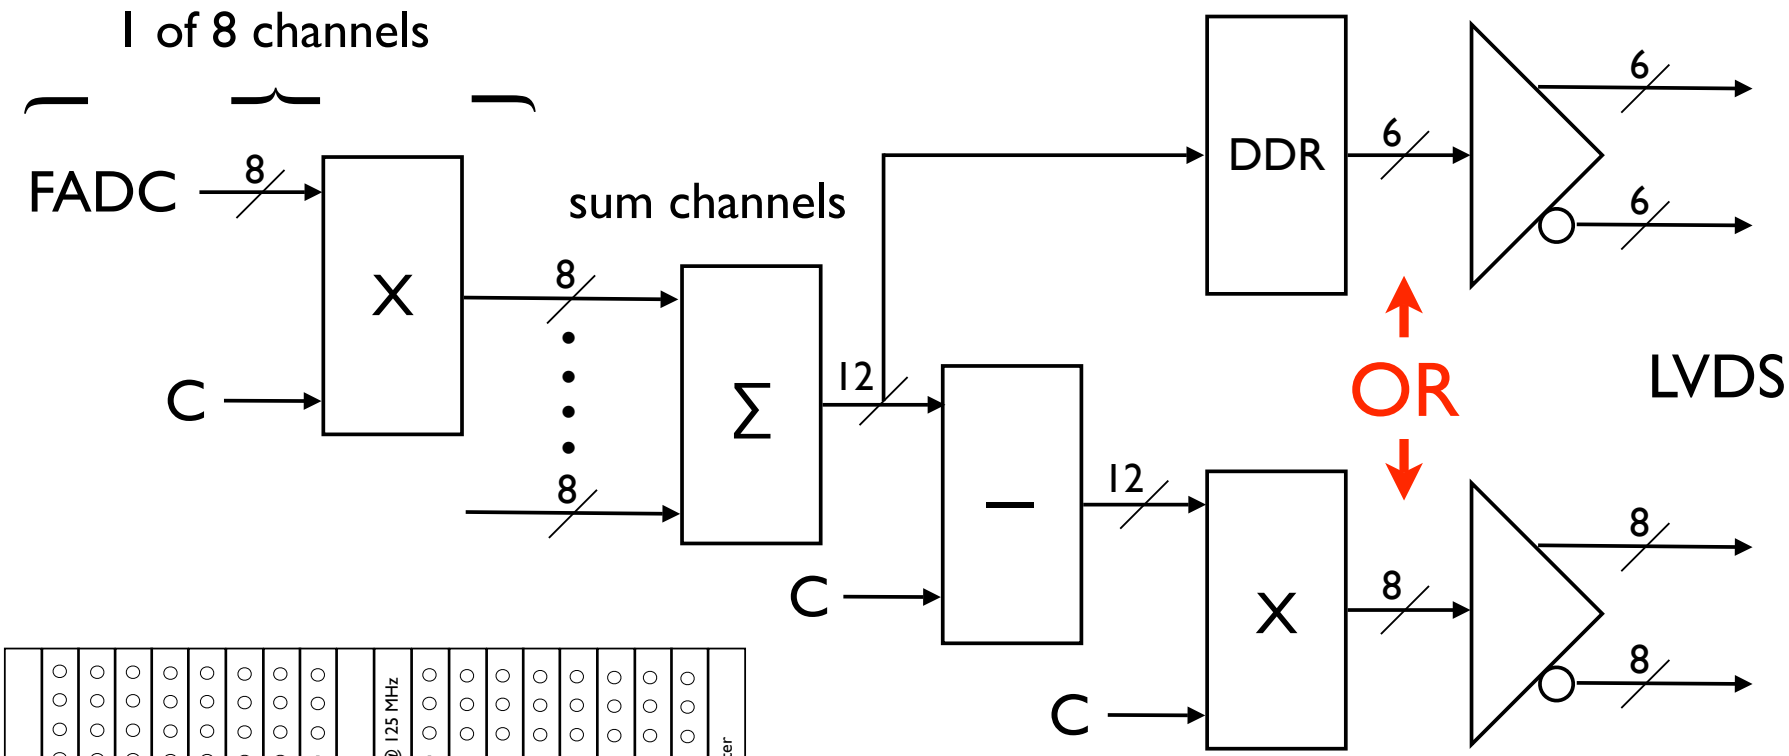

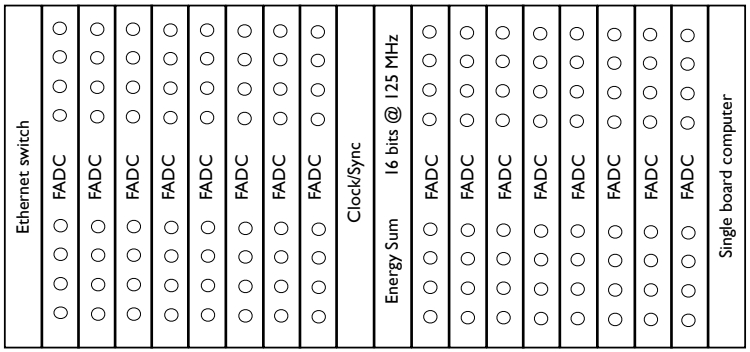

PICMG 2.17
or VXS
Backplane

20 crates FCAL, 8 crates BCAL

Why cPCI?

- Commodity silicon - direct connection to FPGAs, etc.
- IP (Xilinx, opencores.org, etc.)
- Inexpensive shielded 2mm connectors
- Lots of ground & user defined (rear I/O) pins
- “High Availability” features/standards
- Crate management standard (power, fans, modules)
- PICMG: 450 companies

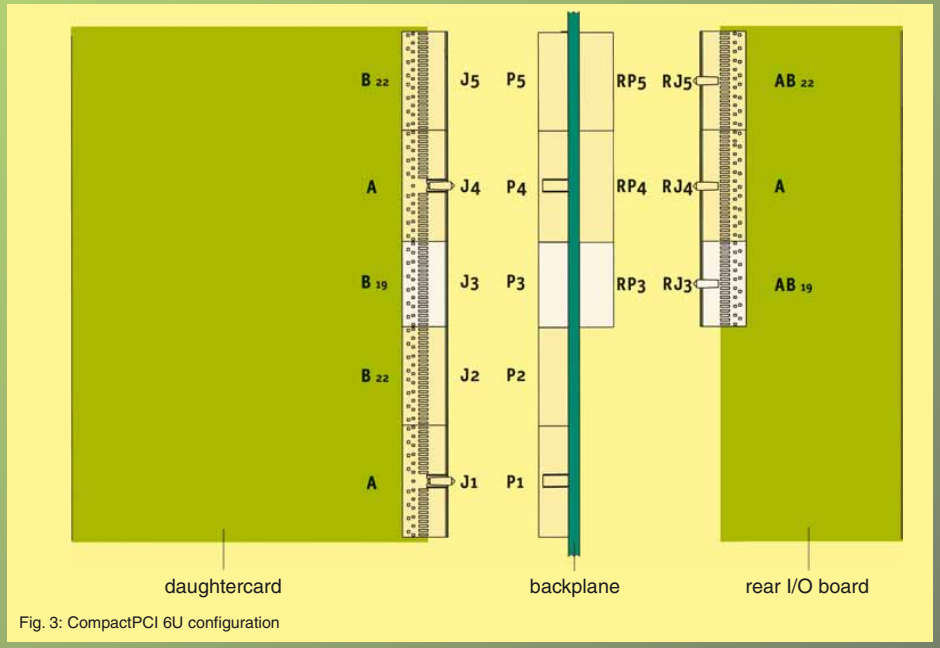

Fig. 3: CompactPCI 6U configuration

Internal Report LBNL-54803

LBNL Report of the Vetting Review of the GRETINA Project

The GRETINA Vetting Review was held on Nov. 4 & 5, 2003

Date of this Report: Jan. 12, 2004

with Addendum of Mar. 29, 2004

J. H. Bercovitz, F. S. Bieser, R. C. Jared, V. P. Karpenko, S. R. Klein, K. T. Lesko, J. E. Rasson,
H. G. Ritter, K. E. Robinson (chair), C. E. Tull, R. Wells, H. H. Wieman
E. O. Lawrence Berkeley National Laboratory

Review Scope and Charge

GRETINA is a gamma-ray detector array capable of reconstructing the energy and spatial positions of gamma-ray interactions within the germanium crystals. It will be used to study the structure and stability of nuclei under various conditions. The new capabilities provided by gamma-ray tracking will give large gains in sensitivity for a large number of experiments, particularly those aimed at nuclei far from beta stability.

A proposal for GRETINA was submitted to DOE in June 2003. It presented the scientific case, the readiness of technical development, the design, the suggested management organizations, and a proposed cost and schedule. The GRETINA proposal received its CD0 approval in August 2003. The CD-1 review will be held on December 3 and 4, 2003, and will be handled by the DOE-N.

The current trend in electronics is away from VME and toward compact PCI (cPCI). There is a much wider selection of cPCI CPU boards than VME based boards, and Cypress no longer makes VME interface chips. The committee suggests that the collaboration consider using cPCI for the next version of the DSP board (and trigger boards).

NIM A 498 (2003) 430-442:

Data Acquisition for DCBA

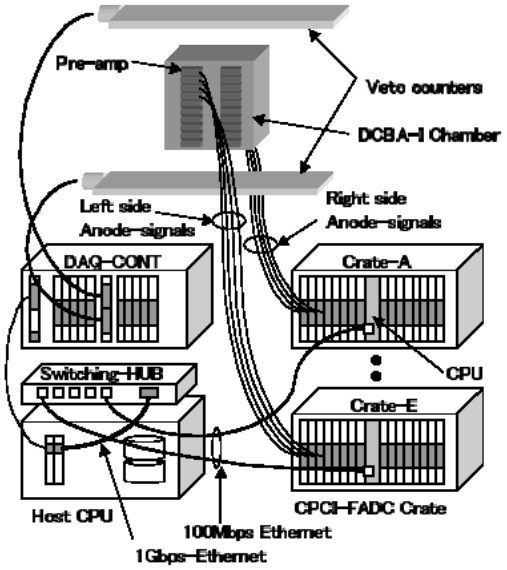

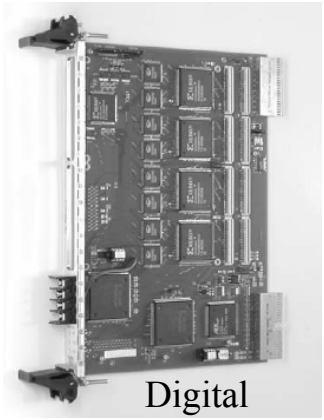

Digital

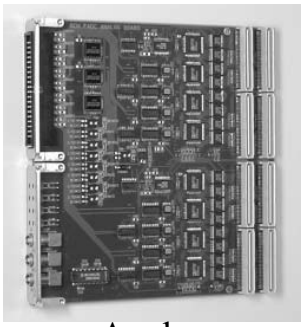

Analog

2 GS/s CompactPCI Digitizer

$8K

CompactPCI™

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.

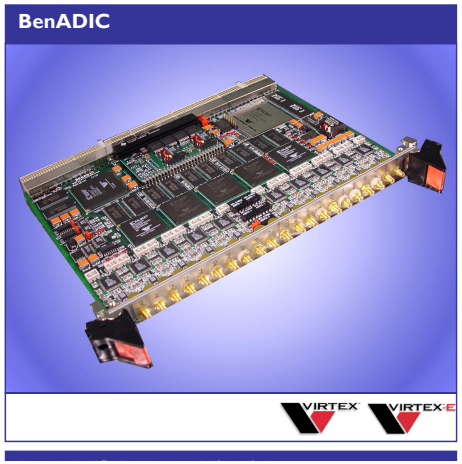

UltraFast

STRATEGIC TEST

UC.2000 - 8 bit transient recorder up to 200 MS/s

- **CompactPCI 6U format**
- **Up to 200 MS/s on 2 channels**
- **Up to 100 MS/s on 4 channels**
- **Simultaneously sampling on all channels**
- **7 input ranges: ±50 mV up to ± 5 V**
- **Up to 512 MSample memory**
- **FIFO mode for slower samplerates**
- **Window and pulsewidth trigger**
- **Input offset up to ±400%**
- **Mutil-card synchronization possible**
- **Windows program SBench 5.x included**

$7.3K

BACKPLANE:

Ethernet-Based Switched Fabric

EtherPlane™

APW Electronic Solutions designs and manufactures a line of high speed CompactPCI Packet Switching Backplanes (cPSB) in compliance to the PICMG 2.16 standard. EtherPlane™, APW's Packet Switched Backplane product line, is now enhanced to include the PICMG 2.16 standardized embedded Ethernet routing for next-generation high-speed packet switch applications. cPSB is a major extension to the CompactPCI Specification with significant potential for embedded system integrators, telecoms suppliers (Voice over IP is an ideal application) and, indeed, anyone who wants to take advantage of benefits of using IP based communications, either with or without the cPCI Bus.

At present, the CompactPCI specification enables single conversations operating at 66 MHz over a 64-bit data bus, a bandwidth of approximately 4 Gbits/s transfer rate, albeit limited to only five slots on a contiguous backplane. The PICMG 2.16 specification allows two switching fabrics, each supporting 20 simultaneous conversations at 2 Gbits/s to provide a 40 Gbits/s transfer rate. Developers can also create "virtual backplanes," expanding to any number of CompactPCI (or non-CompactPCI) systems, by running either fiber or CAT5 Ethernet cables to external connections that extend the packet-switched bus.

CompactPCI/PSB (or cPSB) significantly improves performance, scalability and reliability of CompactPCI while preserving its mechanical, power, hot-swap attributes and H.110 capabilities. System integrators can mix the system components with legacy units relying on the CompactPCI bus and can create interaction within the same chassis. With sub-systems built around legacy CompactPCI elements, system capabilities can evolve organically and gradually onto the CompactPCI/PSB framework. Overall, an extremely significant upgrade, one that maintains backwards compatibility and is achievable without significant engineering investment.

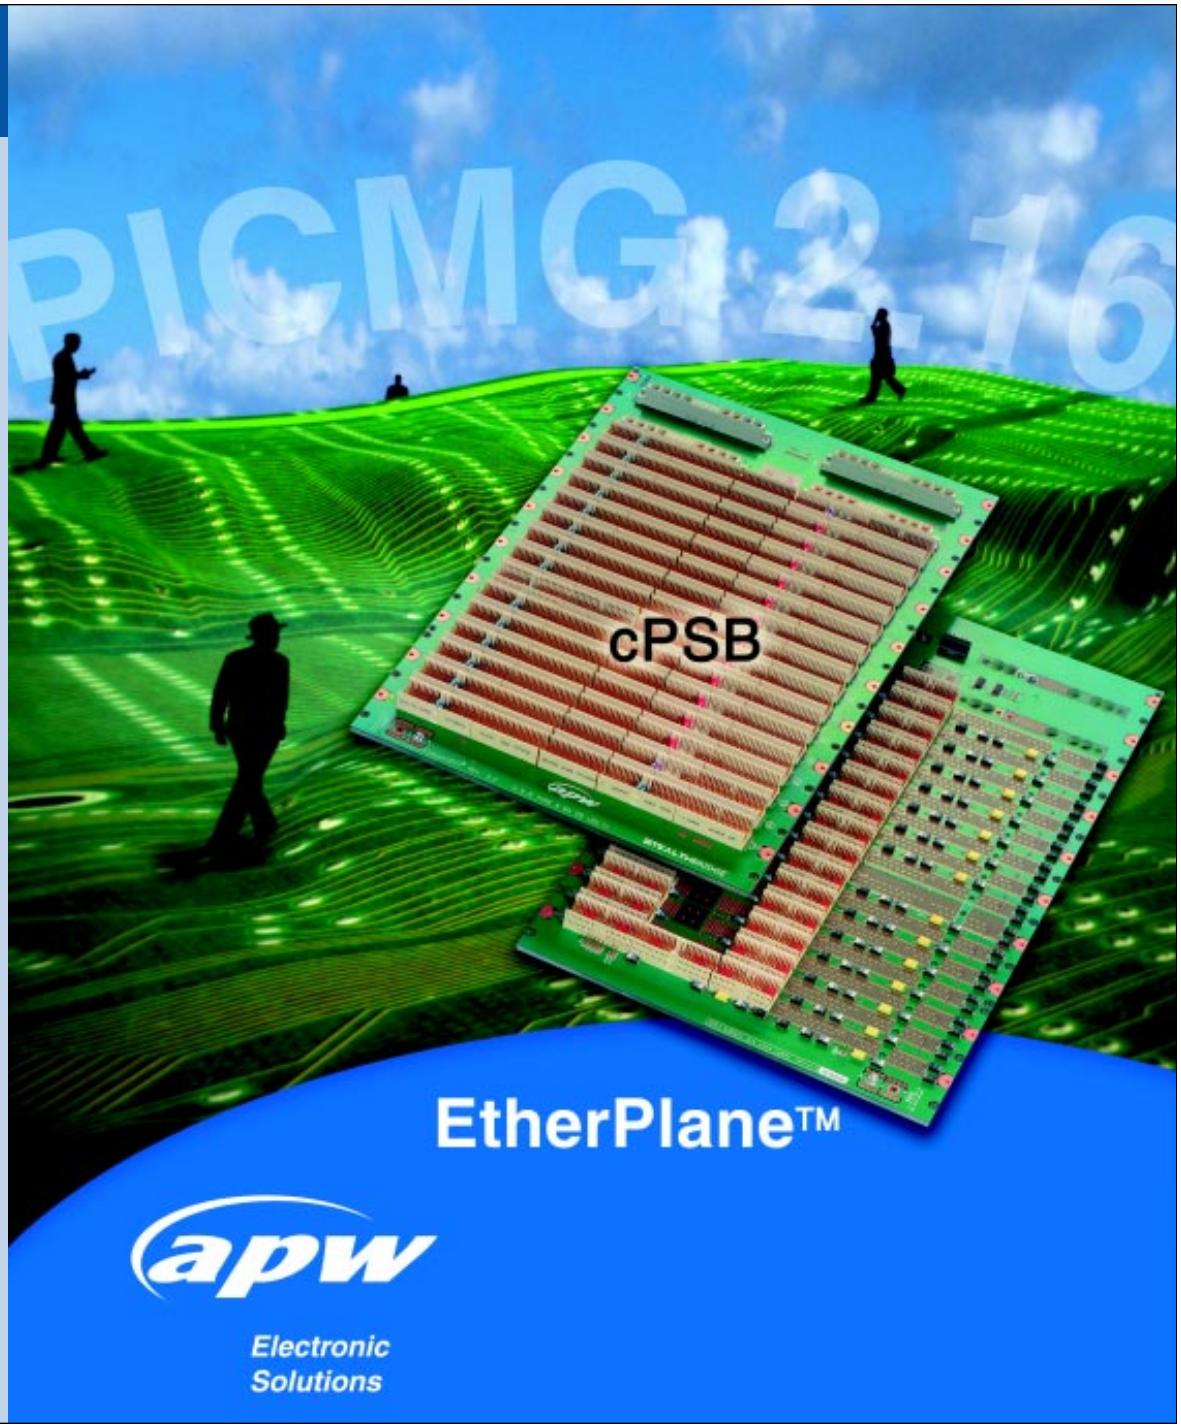

CSB4240

PICMG® 2.16 PSB Ethernet Switch

$1,539

CompactPCI

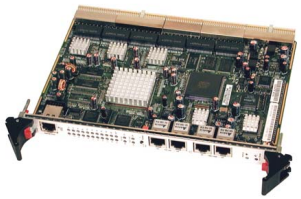

Product Features

6U x 4HP Managed Layer 2/3 Switch

Full Wire Speed on All Ports

Two 100/1000 Ethernet Front Panel RJ-45 Ports

Two 10/100 Ethernet Front Panel RJ-45 Ports

- **XL-ES24**
- **CompactPCI Fast Ethernet switch**

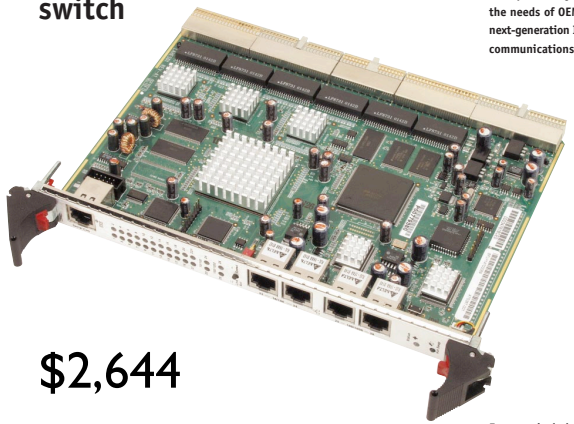

$2,644

- **High-performance, maximum density and ultra-fast connection speeds**

Kontron's XL-ES24 CompactPCI Fast Ethernet switch is a 6U, high-performance, managed layer 2/3 switch which provides twenty-two 10/100 Mbps Ethernet ports for PICMG 2.16 compliance in-chassis switching duties. The XL-ES24 also comes complete with two front panel 10/100 RJ-45 ports and two Gigabit Ethernet ports for fast connection speeds and flexibility.

The XL-ES24 offers maximum density and improves reliability while minimizing external wiring. Console access is enabled through an RS-232 serial cable to configure SNMP, Telnet, CLI and RMON management functionalities; and the XL-ES24 features an easy-to-use browser/Web-based management console – routing and switching at full wire speed utilizing non-blocking architecture.

- **As part of the XL family of high performance CompactPCI building blocks, the XL-ES24 is at once highly integrated and scalable; and specifically designed to meet the needs of OEMs developing next-generation Internet and voice communications networks.**

Features include:

- **Layer 2/3 switching functions**
- **PICMG 2.16 compliant for in-chassis switching**
- **Twenty two, 10/100 Fast Ethernet ports to midplane connectors**
- **Two 10/100 RJ-45 ports**
- **Two 100/1000 RJ-45 ports**
- **Full wire speed on all ports**
- **VLAN configuration distribution**
- **Hot-swappable with LED indication**
- **Status LEDs for Ethernet port link, speed and activity**

kontron

RadiSys

PICMG 2.16 COMPLIANT

DATASHEET

ESM-3100

$3,131

COMPACTPCI SWITCH AND IPMI MANAGER

FEATURE SUMMARY

- Single slot CompactPCI PICMG 2.16 compliant
- Switch and IPMI Manager
- Switch:
  - 24 Link ports fast Ethernet Layer 2 Switch
  - 4 Gigabit Ethernet up-links, 2 optical and 2 copper (with RTM)
  - Latency < 4 us
  - Wire-speed switching
  - Advanced Layer 2 protocol support: VLAN, GMRP, Spanning Tree, Link Aggregation
  - DHCP, TFTP, FTP support
  - Remote management capability via SNMP
  - Command Line Interface to setup and program switch
  - Flash Memory to support custom configuration
  - Optional Rear Transition Modules (RTMs)– single slot wide
  - IPMI (PICMG 2.9) enabled
  - Front panel activity indication
- Integrated IPMI Manager:
  - Dynamic configuration of the system
  - SNMP management interface to hardware platform
  - Power and temperature monitoring
  - Distributed Hot-swap monitoring & control
  - Control and monitoring of 12 General Purpose signals
  - Remote resetting, power cycling

ESM-3100

ESM-3100, is single-slot CompactPCI PICMG 2.16 compliant Layer 2 Switch featuring 24 Fast Ethernet link ports, 4 Gigabit up-links, and integrated platform management. The ESM-3100 is designed to be a companion product to the RadiSys CP50 chassis, providing the switching capabilities for PICMG 2.16 compliant chassis as well as IPMI (PICMG 2.9) management support.

24 PORT FAST ETHERNET PLUS 4 PORT GIGABIT ETHERNET PERFORMANCE

Systems with multiple I/O boards are able to process large amount of data and since almost all of the data in a Compact PCI Switched Backplane (CSBP) is routed through the switch it is essential for the switch to have high-bandwidth connections to pipe the data in and out. For this reason the ESM-3100 is equipped with four Gigabit ports, two optical Ethernet ports on the front panel and two Gigabit copper ports at the rear of the switch on the Rear Transition Module (RTM).

All intra-chassis connections are sup-ported via Fast Ethernet ports. The ESM-3100 can operate as either a stan-dard PICMG 2.16 fabric card, with connections for up to 19 PICMG 2.16 nodes slots, or an extended fabric card with support for up to 24 node slots.

OA&M—OPERATIONAL ADMINISTRATION & MANAGEMENT SUPPORT

The ESM-3100 is remotely manageable via standard SNMP interfaces through a 10 Base-T Ethernet port out the front panel or a 10/100 Base-T Ethernet connection out the rear, through the RTM, or through a virtual interface via the switch ports. In addition to the standard SNMP interfaces, the ESM-3100 is a fully IPMI (PICMG 2.9) compliant blade. This allows the ESM-3100 and other blades in the systems to be controlled and monitored through IPMI commands including resetting of blades remotely, monitoring of temperature, power and logging of other status information made available to the IPMI interface.

IPMI supported commands are also accessible through SNMP via the SNMP agent that runs on the ESM-3100.

The ESM-3100 is equipped with up to 32MB of Flash memory to enable the switch to power-up in a custom setup and/or enable the user to load custom code that will be loaded at power-up.

Custom setups are configurable via a CLI (Command Line Interface) that is accessible over RS-232 or Telnet or the Internet through a Web interface. Custom code is downloadable to the Flash via FTP.

http://recycle.lbl.gov/~ldoolitt/icalepcs2003/WE601.pdf

EMBEDDED NETWORKED FRONT ENDS - BEYOND THE CRATE*

Lawrence R. Doolittle, LBNL, Berkeley, CA 94720, USA

The current generation of Field Programmable Gate Arrays (FPGAs) is an enabling technology, providing flexible and customizable hard-real-time interfacing at the downloadable firmware level, instead of the connector level. By moving in the direction of a system-on-a-chip, improvements are seen in parts counts, reliability, power dissipation, and latency.

This paper will discuss the current state-of-the-art in embedded, networked front end controllers, and gauge the direction of and prospects for future development.

While analog electronics has not shrunk as dramatically as digital, it has proved possible in many cases to simplify the analog signal path by pushing functionality into the digital domain[4]. This is an important step in bringing down the total hardware complexity, since the digital processing involves no additional chips.

The flexibility and end-to-end integration of an FPGA-based SOC make it plausible to use Ethernet with a hard real time mind-set that is inconceivable using a CPU and a conventional MAC. Frame preamble and header information can be sent down the wire while results are still being acquired from the hardware.

FPGAs are an enabling technology. Their reconfigurability is an essential feature, allowing bugs to be fixed and features to be added to the hardware at a later date. This flexibility comes with a hardware price: some means of "booting" or "configuring" the FPGA must be included, and (to avoid losing the very feature that is so attractive) a mechanism must be included to make that configuration remotely updatable. When a conventional networked com-

http://epaper.kek.jp/ica01/papers/WEAP023.pdf

8th International Conference on Accelerator & Large Experimental Physics Control Systems, 2001, San Jose, California

WEAP023
cs.DC/0111033

MODERNISING THE ESRF CONTROL SYSTEM WITH GNU/LINUX

A.Götz, A.Homs, B.Regad, M.Perez, P.Mäkijärvi, W-D.Klotz
ESRF, 6 rue Jules Horowitz, Grenoble 38043, FRANCE

The result of this technology survey was 100 MHz Ethernet, VME (for the existing hardware), CompactPCI (cPCI) and PCI for new hardware, Linux as the main frontend op-

3 LINUX/M68K + VME

The ESRF has over 200 VME crates installed. This represents an investment of millions of Euros as well as many tens of years of work in hardware and software development. Any modernization project must take this investment into account. The modernization foresees two ways to do this: using the Motorola CPU's (MVME-162) to run GNU/Linux directly, or replacing the CPU with a bus extender which allows the VME bus to be controlled from PC running Linux/x86. This section describes the first option.

4 LINUX/X86 + BUS EXTENDERS

The modernization project of the instrument control at the ESRF using GNU/Linux supports two main hardware platforms: PCI/cPCI and VME. The former provides access to the most recent interface boards developed for a highly demanding market, and hence, with better performance/price ratios. The latter is needed for a gradual transition between the current VME instrumentation and the PCI technology. VME boards can be controlled from a Motorola MVME CPU or from a PC through a PCI/VME bus extender, both running GNU/Linux as OS.

High Speed Data Acquisition and Trigger

Walter F.J. Müller
GSI

http://www.veccal.ernet.in/~icpaqgp/WFJMueller.pdf

DAQ Architecture - Future

Self-triggered Data Push Architecture

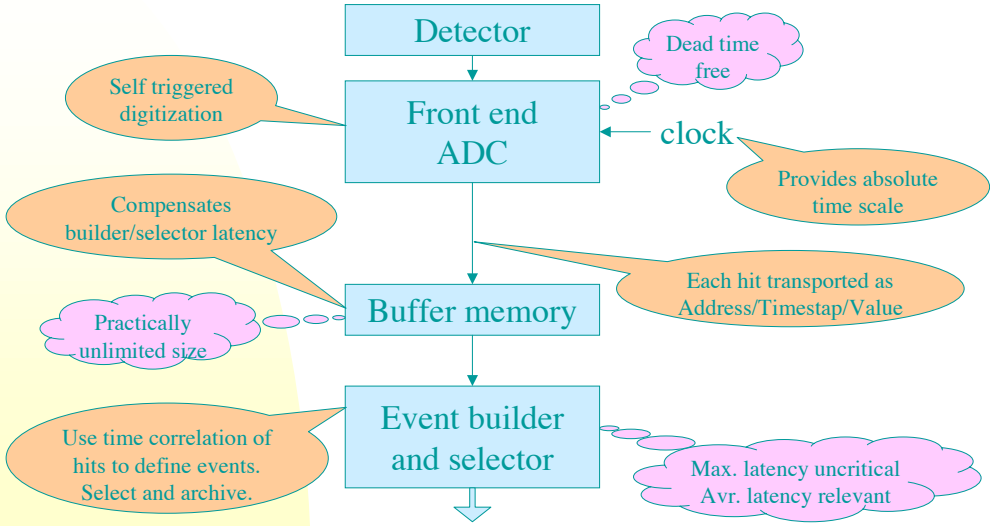

Planned Experiments with SDPA

- AGATA (Advanced Gamma Tracking Array — Completion ~2008)
  - 190 Ge detectors – 6780 channels
  - 300 kHz events/sec @ M=30 (DC beam)
  - ~ 1 GB/sec into reconstruction farm
- BTeV (Charm & Beauty Decays at FNAL — Completion ~2007)
  - $2.2*10^7$ pixel + RICH + ECAL + ….
  - 7.6 MHz bunch crossing (Bunched beam)
  - ~ 1 TB/sec into L1 buffer memories
  - L1 trigger on displaced vertices

Crates and Backplanes

10 Gbps SERDES in CMOS

- Trend: use serial point-to-point links
  - Parallel `shared media' busses obsolete
  - Look for serial backplane fabrics
- Backplanes: What's available today/tomorrow ?
  - PICMG 2.16: C-PCI + dual 1G Ether star (available)
  - PICMG 2.17: C-PCI + 4*622 Mbps star (available)
  - PICMG 2.20: C-PCI + 2.5 Gbps mesh (2.16+2.20 announced)
  - VITA 41 (VXS): VME + 4*10 Gbps dual star (Infiniband over P0 conn.)
- What's in the pipe ?
  - ATCA (Advanced Telecommunications Computing Architecture)
    - Base Interface: dual 1G Ethernet star
    - Fabric Interface: 8*10 Gbps star or mesh