

L2Alpha Processor

Robert D. Martin

University of Illinois at Chicago


6 February 1999

L2Alpha Requirements

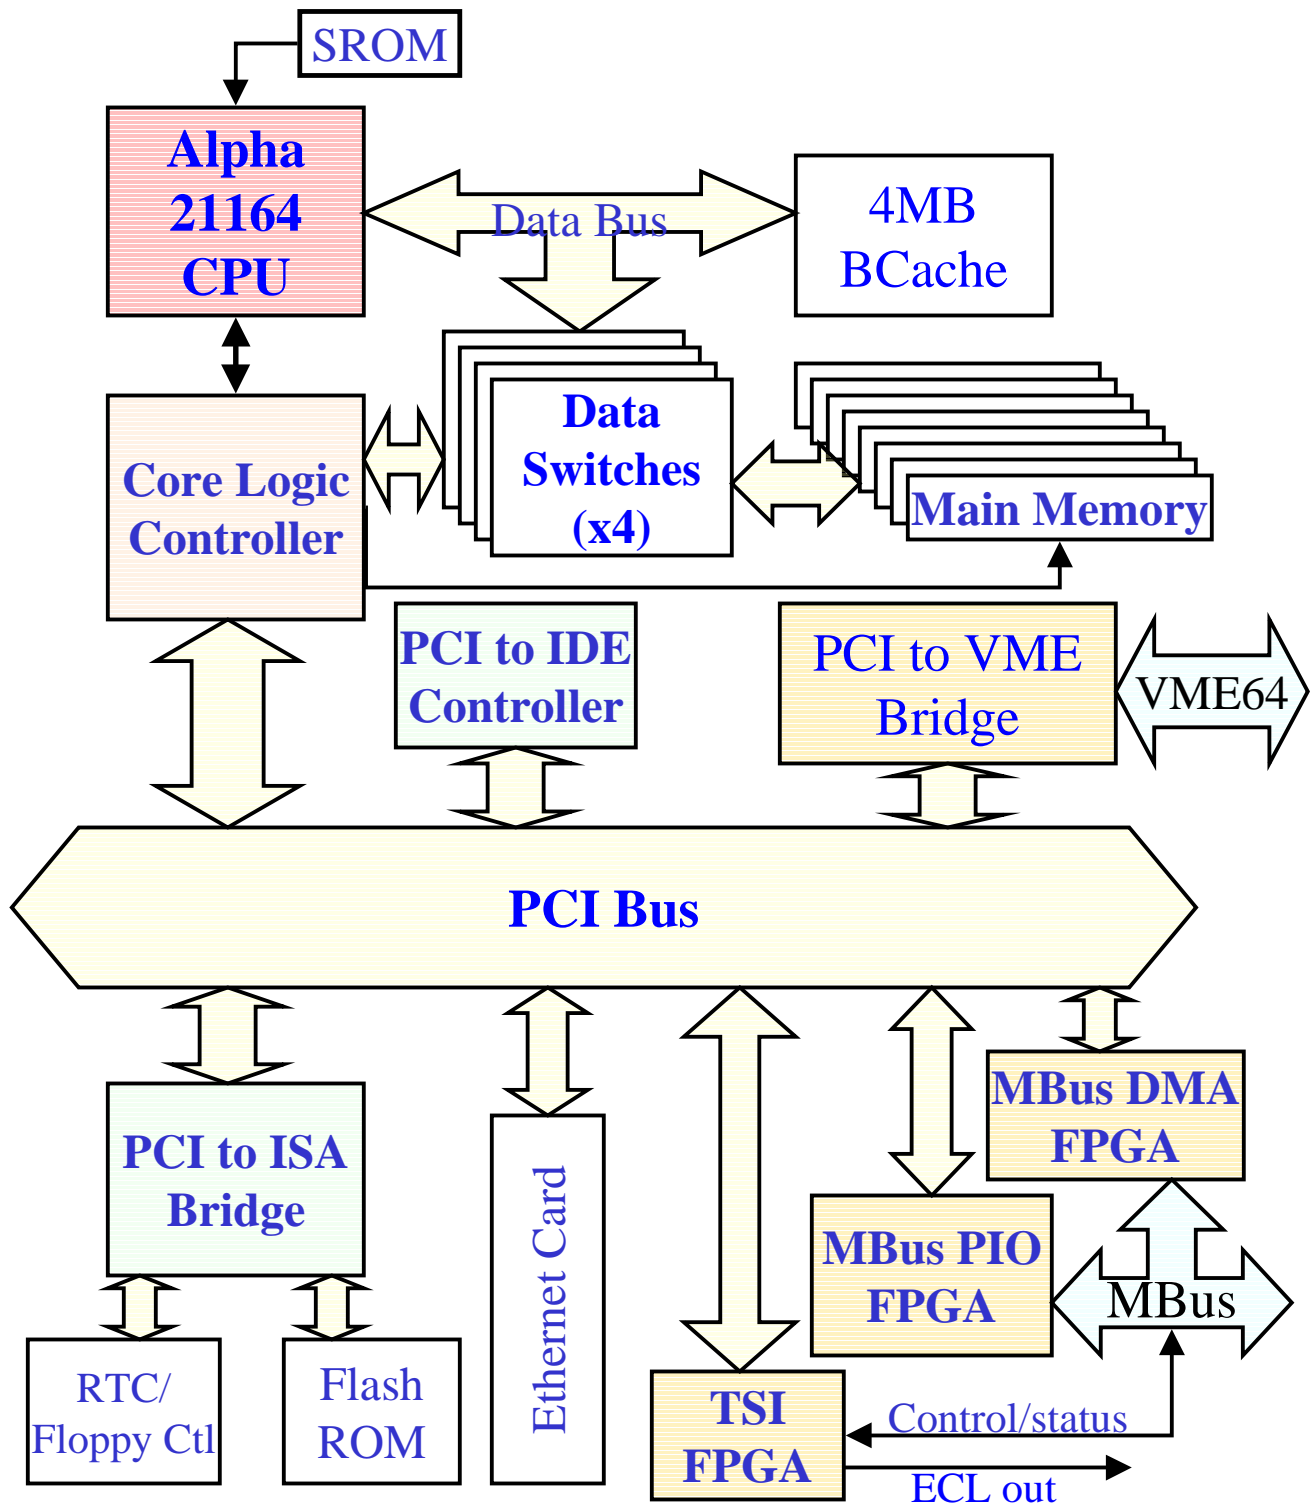
- **Fast CPU**
 - Event processing within 100 μ s
- **VME Interface**
 - Communication with TCC
 - Upload of Events to L3 via VBD
 - VME interrupts enabled
- **MBus Interface**
 - High speed for input data
 - Low latency for Administrator--
Worker communication
- **Buffers**
 - Input: at least 16 events
 - Output: at least 8 events

L2Alpha Design

- Joint CDF and D0 project
- Designed by Myron Campbell's group at U of M
- Based on DEC PC164 motherborad
- Tundra UniverseII for VME
- 3 MBus FPGAs



L2Alpha Block Diagram



L2Alpha: PC164

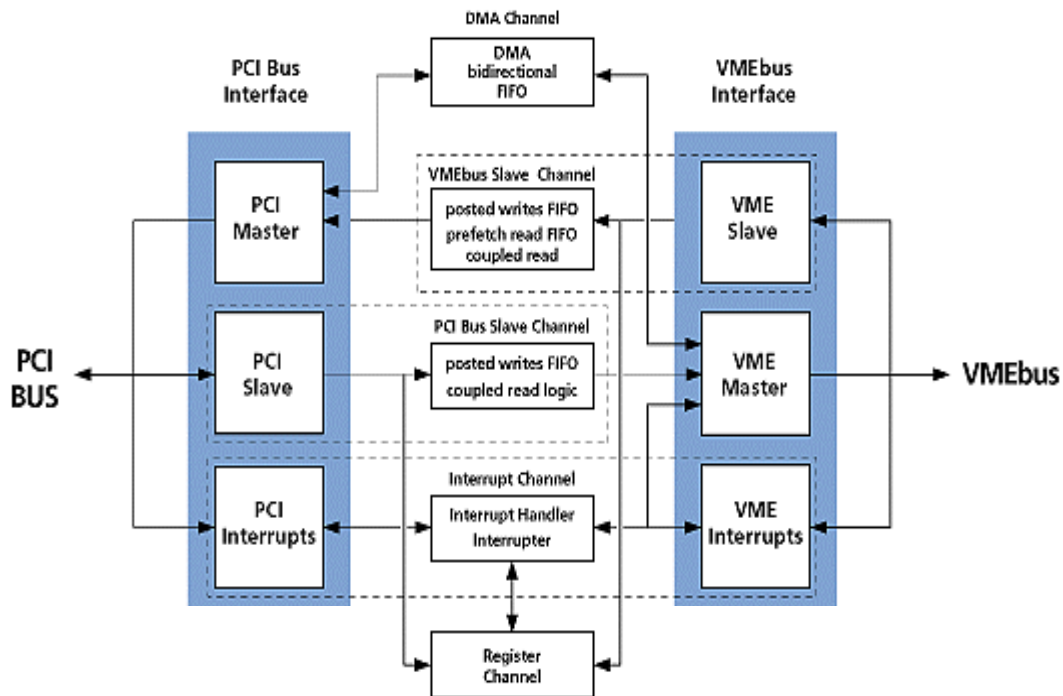
- 500 MHz Alpha 21164 CPU
 - 4 pipelines: 2 int, 2 fp
 - 3 levels of cache
 - separate instruction and data L1 cache (Icache and Dcache)
 - unified 3-way set associative L2 cache (Scache)
 - Off-chip direct mapped L3 cache
- 4 MB L3 Cache (diff from PC164)
- 256-bit wide path from CPU to SIMM Memory
- memory mapping compromise between fast DMA and fast processor access

L2Alpha: PC164 (cont)

- 2 levels of firmware
 - SROM: hardware configuration
 - Flash ROM
 - programmable on the fly
 - OS/task specific configuration
- 64-bit PCI bus (33 MHz)
- 2 PCI expansion slots
- Support for floppy and IDE hard drives
- Real Time Clock with periodic interrupt capabilities
- 3 serial ports: COM1/2 and SROM

L2Alpha: VME

- UniverseII provides VME interface
- VME master or slave
 - 8 address windows each
 - Windows configurable over VME
- Mapping for all 7 VME interrupts
- 1 PCI interrupt



L2Alpha: MBus DMA

- MBus DMA address
 - mbus_ad<7:0> set DMA address
 - mbus_ad<31:8> must be zero
 - 256 possible DMA addresses
- DMA uses 2 PLDs
 - FPGA for DMA addr to PCI map
 - CPLD for xfer to main memory
- MBus transfer and PCI transfer decoupled by 1K deep FIFOs
- When FIFO not empty
 - translation made between MB address and PCI address
 - 64 bit PCI DMA transaction to main memory until FIFO empty

L2Alpha: MBus DMA

- Translation Buffers for each MBus Address
 - Starting PCI address in translation buffer configurable for each MBus broadcast address
 - Each 64 bit PCI transfer increments address in translation buffer
- Above features can be used to determine xfer length for error detection
- Open Collector line pulled while data in FIFO

L2Alpha: MAGICFPGA

- MAGICFPGA does programmed I/O (PIO) transfers on MBus
- 2 PCI to MBus windows
 - 1 window must xfer 128 bits
 - 1 window can xfer < 128 bits
- 1 MBus to PCI window
- 128 bits are always transmitted to or received from MBus
- Only path for non-DMA MBus communication between MBus devices

L2Alpha: TSI FPGA

- “Catch-all” FPGA
 - Mbus control and status lines
 - MOD_DONE<21:7>
 - AP_FIFO_EMPTY
 - START_LOAD
 - etc
 - Can pull PCI interrupt -- will be used for new event arrival
 - 32 ECL scalar channels driven by TSI -- diff between CDF/D0

I/O Paths

- Ethernet
 - slow
 - only appropriate for d/l of executable in running system
 - can be used for remote debugging
- VME
 - output to L3
 - monitor output to TCC
 - run-time parameter d/l
- MBus
 - Input data
 - Output to L2Global
 - Interprocessor communication
 - Communication w/ MBT

Interrupts

- Sources
 - Real time clock
 - useful for monitoring/debugging
 - minimize (few Hz)
 - MBus (1 int)
 - originating on TSI-FPGA
 - Will use to signal when new event has been moved from DMA FIFOs to memory
 - VME
 - 7 VIRQ go to 1 PCI, mapping in Universe
 - Will be used for asynchronous SCL and TCC information

Software

- Software Developers Kit (SDK) from Digital
- “Debug Monitor” Firmware (DBM)
 - Allows board to run w/o an OS
 - Access to full physical memory
 - Dynamic mem alloc should be avoided
 - Executables can be d/l’ed over ethernet using BOOTP/TFTP
- Remote debugging over ethernet from Digital UNIX using ladebug.
- Can modify DBM to load trigger exe on startup

Downloading and Verification

- Program download over Ethernet.
 - < 4MB program over 0.5 MB/s link
 - < 10 sec per node
 - can run immediately after d/l
- Parameter d/l by TCC over VME
 - << 1 MB over 2 MB/s link
 - more info given by R. Moore
- Where possible, lookup tables should be calculated not downloaded
- Can run *in situ* test programs for software verification

Project Information

- Boards designed by Myron Campbell's group at U of M
 - Stephen Miller
 - Zhihui Huang
- Prototype MBus DMA engine and Universe expansion cards built and tested prior to integrating those designs into L2Alpha
- Fab/Stuffing by ADCO and their sub-contractors
- D0 will build roughly 3x more L2Alphas than CDF

Prototypes

- 7 Prototypes delivered to UofM in December -- 5 go to D0 institutions
- In process of design verification
 - Most of PC164 functionality tested and verified (some corrections made)
 - Working w/ Tundra engineers to understand UniverseII problem
 - MBus FPGAs not tested yet
 - Ignoring obvious manufacturing problems until have working design
 - should be done by March workshop

After the Prototypes

- Some design mods need to be done
 - fixes to problems found w/prototypes
 - additional functionality
 - add drivers to make some MBus lines R/W instead of Read Only
 - other possible mods seen as necessary once start operating boards
- “Best Guess” of 12 weeks for production fab/stuffing

Production Debugging

- Design Verification
 - use first CDF or D0 version (both?) to verify production design
 - Work with UofM/CDF
- Board debugging
 - Will be done at UIC
 - HEP group + department EE

Longterm Support

- Have enough spares on hand (on shelf or in test crates) to only need to repair boards a couple times per year
- 3am replacement not 3am repair
- Given the resources, repair will be done by the UIC

Longterm Support Issues

- L2Alpha is based on a commercial computer
- Most component parts have life cycle of computer industry not HEP
- Need to purchase spare parts prior to being discontinued
 - non-EDO 72pin SIMMs may already be a problem!
 - May need to corner market of PCI to ISA bridge used on prototype/PC164
 - **Will only get worse!**