


The Compute Node / ATCA System

Björn Spruck for the Giessen Group

II. Physikalische Institut, JLU Gießen

20.10.2011

Online Tracking EVO Meeting



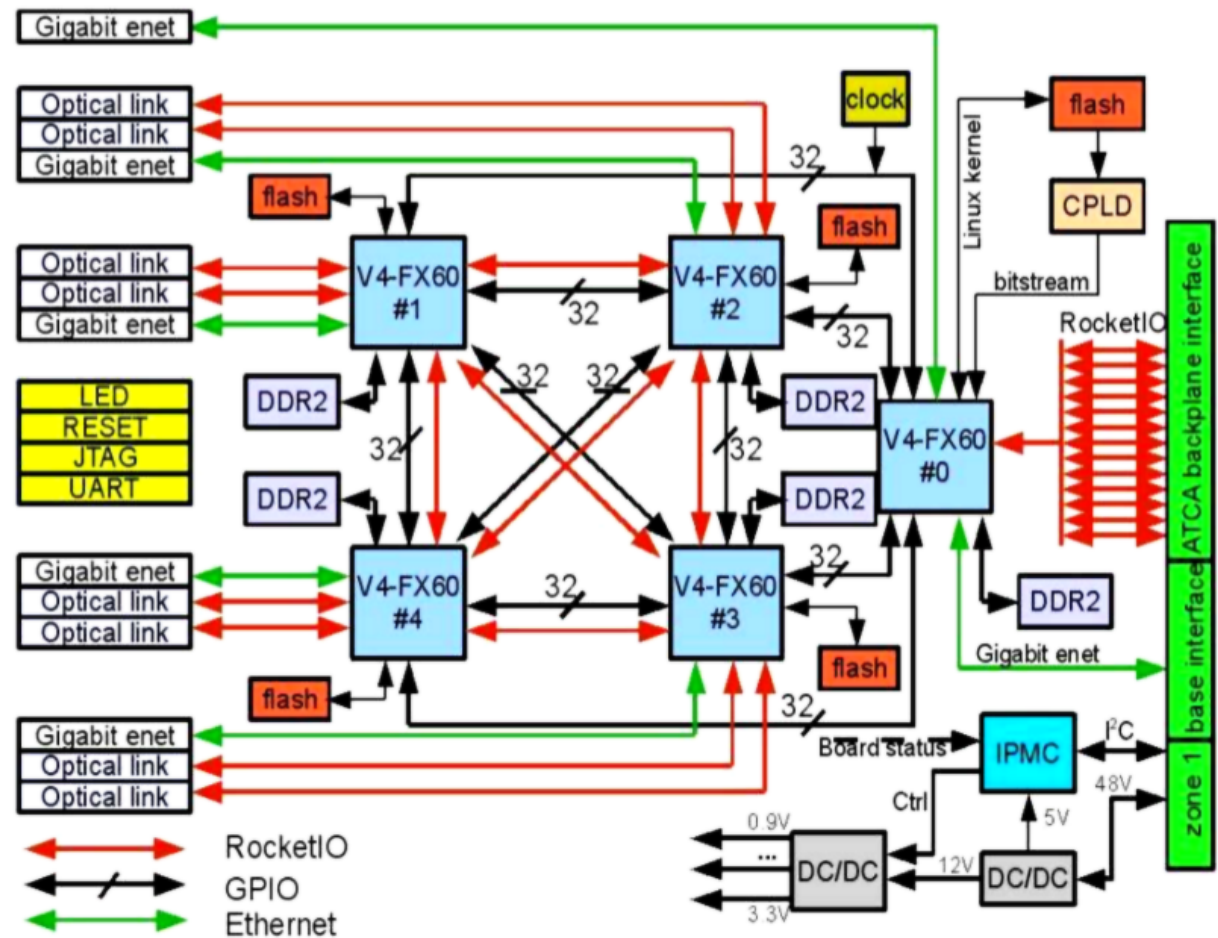
II. Physikalisches
Institut




JUSTUS-LIEBIG-
 UNIVERSITÄT
GIESSEN

bjoern.spruck@physik.uni-giessen.de


The Compute Node (CN)

- Developed in close cooperation with Trigger Lab of IHEP Beijing, China
- High Performance Computing
 - 5 Virtex-4 FX60 FPGA
 - (upgrade: Virtex 5)
 - 5*2GB DDR2 RAM
 - (upgrade: 4GB)
 - interconnected by RocketIO
- ~32Gbps Bandwidth
 - 8 Optical Link (3Gbps each)
 - (upgrade: 6.5Gbps)
 - 5x Gigabit Ethernet
- 13x RocketIO to backplane (full mesh)
- 2 embedded PowerPC in each FPGA for slow control
- ATCA compliant





- IPMI remote control
- Small piggy-back board
- Functions: power on/off, power negotiation, health monitoring, board reset, bitstream selection



3 kW Power Supply


Full Mesh Backplane

Primary Shelf Manager

Backup Shelf Manager

Fan Trays

- Two prototype boards, one in Gießen
 - Daughterboard can run standalone!
- Motherboard close to PCB production



- Operating System:
 - PowerPC embedded linux, kernel 2.6.37 (from 2011)
 - Two setups
 - Standalone from initramfs
 - Complete system by NFS if available; including gcc, ssh, ...
 - Software TCP/IP (slow)
- Full memory is accessible by PowerPC, but not controlled by Linux
 - Full memory and hardware addresses (IP core) accessible from user program
- Memory access from PowerPC and IP Core → cache problems possible!
 - Use hardware registers for control
- **PowerPC controls the IP cores, but does no computing itself.**
 - **Finally, no data transferred by PowerPC, only “slow” control**

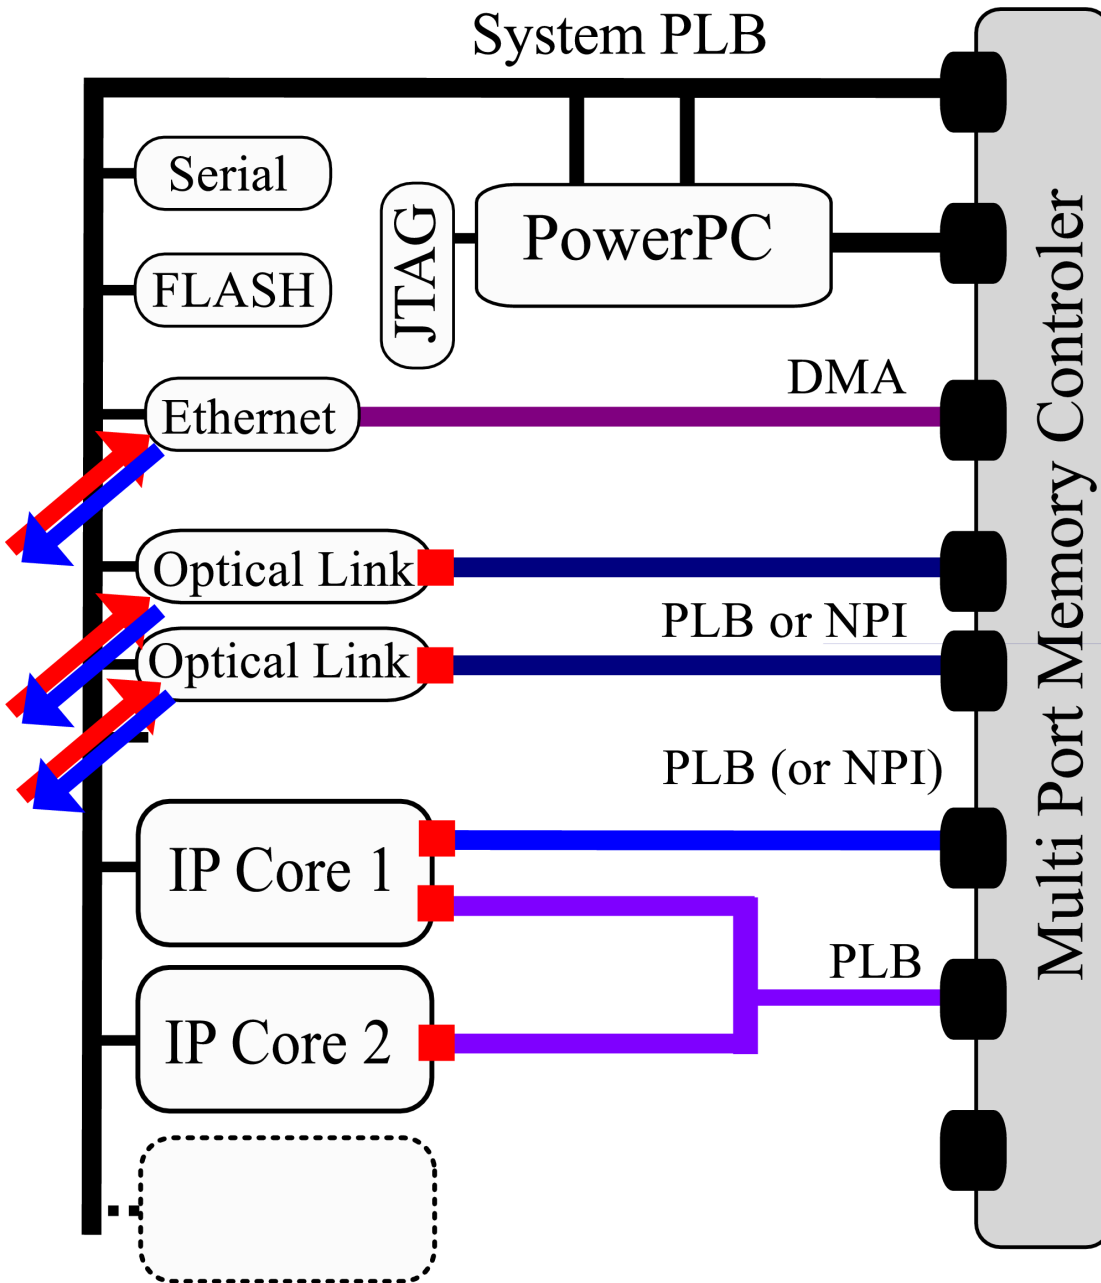
Linux on Virtex4 and 5, Example "Firmware"

The screenshot displays a terminal window with a dark background and white text. The terminal output is divided into several sections:

- Network Traffic:** Multiple lines of data showing network statistics such as "C0: 0 C1: 999957 C2: 0 C3: 0", "100 % buffers empty", and "Last Offs=08EF3000, Last Id=996706, HLT 0 Evt 0".
- System Information:** A section titled "root@localhost:~/src#" showing the output of the command "cat /proc/cpuinfo". It lists details for a Virtex-5 FXT processor, including clock speed (400.000000MHz), revision (25.18), bogomips (800.00), timebase (400000000), platform (Xilinx Virtex440), model (testing), and memory (768 MB).
- SSH Connection:** A section titled "root@cn2:~#" showing the output of "cat /proc/cpuinfo" after an SSH connection to a remote host. It lists details for a Virtex-4 FX processor, including clock speed (300.000000MHz), revision (20.112), bogomips (600.00), timebase (300000000), platform (Xilinx Virtex), model (testing), and memory (512 MB).

Two pink arrows point from the terminal output to the system information sections. In the background, a window titled "Terminal" shows a network diagram with two nodes labeled "xc4vfx60" and "load_belle2_..." and "download_belle2_...".

FPGA - Virtex 4/5




Simple data flow:

- Receive by optical link
- Preprocess / write to memory
- Read buffered data
- Process data
- Write to buffer
- Send out by optical link

IPMC

4 Computing FPGAs

Switch FPGA



4 × (Gbit Ethernet + 2 × 2 Gbps Optical)