# ROM Status Update

U. Marconi, INFN Bologna

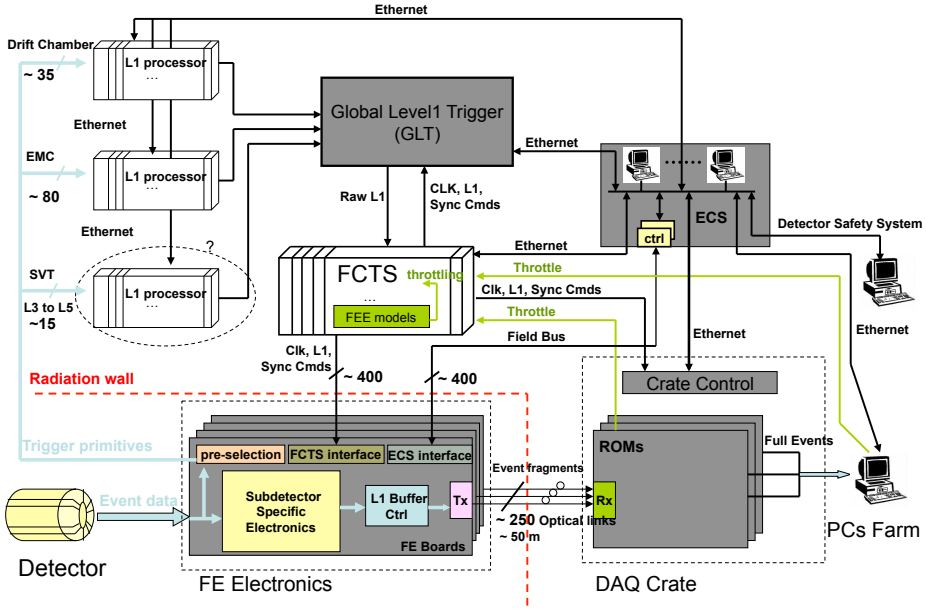


# Assumptions

- Level-0 trigger rates **at $L=10^{36}\ cm^{-2}s^{-1}$** :
  - 50kHz Bhabha, 25kHz beam backgrounds, 25kHz physics+backgrounds
  - 50% headroom desirable
- Baseline: 150kHz rate capability
- Event size: 100 kByte
- Pre-ROM sizes: 500 kByte

# ROM implementation

- FPGA – PC implementation, suitable to work at the baseline expected trigger rate.
- From the FPGA data enters the PC through the PCIe interface.
- We are using a Xlinx Virtex-6 FPGA connectivity kit, equipped with PCI Express (x4 PCI Express Gen2 Endpoint).
- We reached a 14.5 Gb/s transmission rate of a theoretically available bandwidth of 16 Gb/s.
  - Without data processing.
- The PC will forward data to the HLT farm exploiting a 10 GbE NIC board.

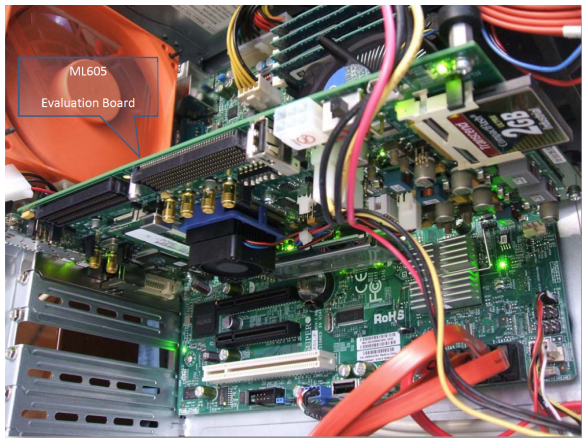


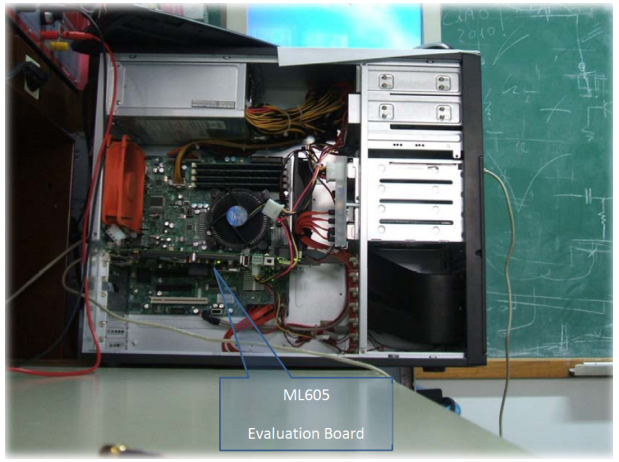

# Exploiting multi-core CPU

- We made preliminary tests distributing event fragments for processing to different CPU cores.
  - The Linux driver writes event fragments at different RAM locations.
  - Processes running on different cores access assigned dedicated RAM banks to fetch data.
- Performance depends on the mother board architecture and CPUs.
  - Different cores share the bus to access the RAM.
- NUMA PC architecture provides full parallelism.
  - NUMA: Each core is equipped with a dedicated bus/RAM.
  - Intel Quick Path Interconnect.

# Rate, ROMs and CPU cores

- Data flux: 150 kHz × 500 kB = 600 Gb/s
- n. ROM: 600 Gb/s / 10 Gb/s ~ 60
- Event fragment size per ROM: 500 KB / 60 ~ 10 KB
- n. of CPU COREs:
  150 kHz × $T_{proc}$ = 150 kHz × 1 ms = 150.
- n. of CPU COREs per PC = 3 ÷ 4
  - Additional 1 or 2 COREs to manage network traffic.
- Minimum $T_{proc}$ when forwarding event fragments as IP packets not doing any processing.

# Higher rates

- What if the rate would rise up to ~500 kHz?
- The flux per ROM becomes
  500 kHz * 10 kB = 40 Gb/s,
  … too much for a single PC… may be.
- Some more ROM's (4 times) will do the job, but it has a cost: each board (an interface board to the FEE and PC) would handle a small event fragment of just ~2.5 KB.
- Doing what on these small fragments?

# Higher rates

- Level-0 trigger rates at $L=10^{36}\ cm^{-2}s^{-1}$ :
  - 50kHz Bhabha,
  - 25kHz physics+backgrounds
  - 25kHz beam backgrounds
  - 50% headroom desirable
- Baseline: **150kHz rate capability**
- Event size: 100 kByte
- Pre-ROM sizes: 500 kByte

- Level-0 trigger rates at $L = 5.\ 10^{36}\ cm^{-2}s^{-1}$ :
  - 250kHz Bhabha,
  - 125kHz physics +backgrounds
  - 25kHz beam backgrounds,
- **Trigger rate > 400 kHz**
- Event size: 100 kByte
- Pre-ROM sizes: 500 kByte

# FPGA implementation

- The optical interface to the FEE is mandatory.
- Data format conversion (from the detector format to Ethernet, the industrial standard) could be performed by a FPGA.
- It requires implementing IP protocol on the FPGA, which will act as a 10 GbE network interface.
- The 10 gigabit Ethernet standard encompasses a number of different physical layer (PHY) standards: 10 GbE optical fiber.

# FPGA 10 GbE

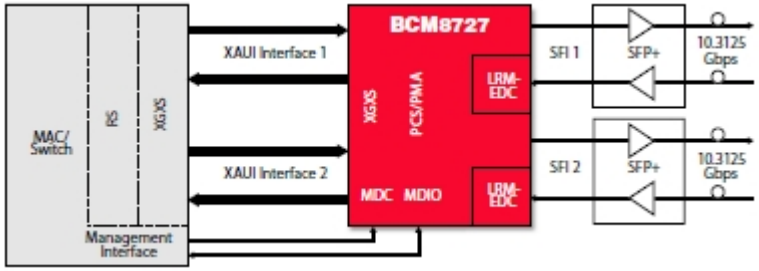
MAC/
Switch
RS
XGXS
XAUI Interface 1
XAUI Interface 2
Management
Interface
BCM8727
XGXS
PCS/PMA
MDC MDIO
LRM-
EDC
LRM-
EDC
SFI 1
SFI 2
SFP+
SFP+
10.3125
Gbps
10.3125
Gbps

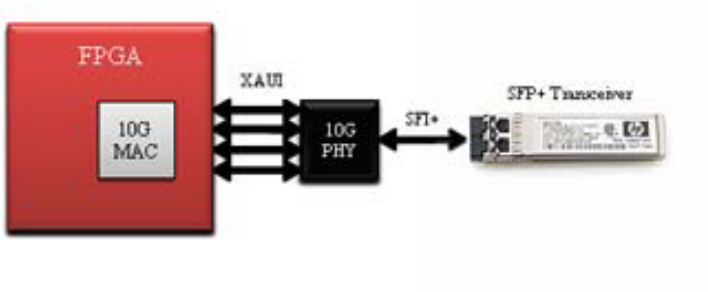
FPGA
10G
MAC
XAUI
10G
PHY
SFI+
SFP+ Transceiver

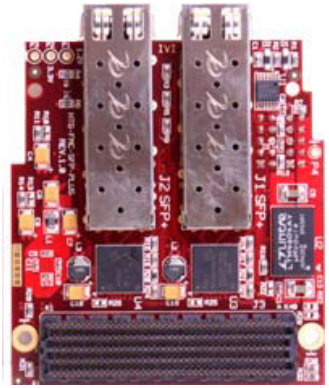

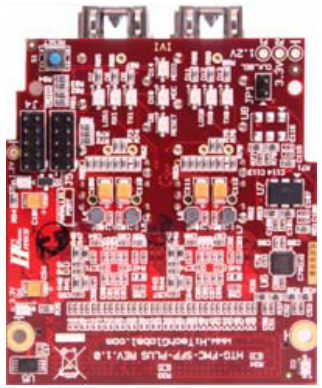

- XAUI to SFP+ conversion
- **XAUI** is a standard for extending the XGMII (10 Gigabit Media Independent Interface) between the MAC and PHY layer of 10 Gigabit Ethernet (10GbE).
- The newest module standard was the enhanced small form-factor pluggable transceiver, generally called SFP+. Based on the small form-factor pluggable transceiver (SFP) and developed by the ANSI T11 fibre channel group, it was smaller still and lower power than XFP. SFP+ became the most popular socket on 10GE systems.[4][5] SFP+ modules do only optical to electrical conversion, no clock and data recovery, putting a higher burden on the host's channel equalization. SFP+ modules share a common physical form factor with legacy SFP modules, allowing higher port density than XFP and the re-use of existing designs for 24 or 48 ports in a 19" rack width blade.
- Optical modules are connected to a host by either a XAUI, XFI or SFI interface. XENPAK, X2, and XPAK modules use XAUI to connect to their hosts. XAUI (XGXS) uses a four-lane data channel and is specified in IEEE 802.3 Clause 48. XFP modules use a XFI interface and SFP+ modules use an SFI interface. XFI and SFI use a single lane data channel and the encoding specified in IEEE 802.3 Clause 49.
- SFP+ modules can further be grouped into two types of host interfaces: linear or limiting. Limiting modules are preferred except when using old fiber infrastructure which requires the use of the linear interface provided by 10GBASE-LRM modules.[6]
- **[edit]**

# Engineering

- Carrier board (CB) to power and control FPGA “format converter units” (FCU).
- Each FCU has n optical input channels and one 10 GbE output (optical) channel.
- Is a completely synchronous system.