A new paradigm using GPUs for fast triggering and pattern matching at the NA62 CERN experiment

G.Collazuol[1,2], G.Lamanna[1,2], M.Sozzi[2,3]

([1] Scuola Normale Superiore of Pisa, [2] INFN sez.Pisa, [3] University of Pisa)

From theory to experiment

■ $K^+ \to \pi^+ \nu\bar{\nu}$ process is unique as extremely accurate and clean probe for the non-trivial flavour structure of physics beyond the Standard Model

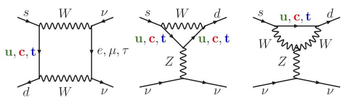

$BR(K^+ \to \pi^+ \nu\nu) = (8.22 \pm 0.84) \cdot 10^{-11}$

■ $K^+ \to \pi^+ \nu\bar{\nu}$ is computed in a clean theoretical environment due to the small contribution by hadronic matrix elements and long distance terms.

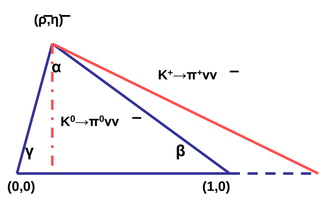

■ FCNC process forbidden at tree level (GIM meccanism)
■ top contribution is dominant in loops: cleanest way to extract V_{td} and to give independent determination of the unitarity triangle

■ **NA62** aims at measuring $K^+ \to \pi^+ \nu\bar{\nu}$ with O(100) SM events, in two years of data taking
■ Very challenging experiment: veto capability on not-interesting decays and high beam intensity
■ tracking of single beam and decay particles, redundant time, momentum, angle measurements, high vetoing power for charged and neutral particles, highly effective particle identification, Ultra high rate capability…

Main backgrounds

$BR(K^+ \to \pi^+ \pi^0)$	$20.7 \cdot 10^{-2}$
$BR(K^+ \to \mu^+ \nu)$	$63.5 \cdot 10^{-2}$
$BR(K^+ \to e^+ \pi^0 \nu)$	$5.1 \cdot 10^{-2}$
$BR(K^+ \to \mu^+ \pi^0 \nu)$	$3.4 \cdot 10^{-2}$

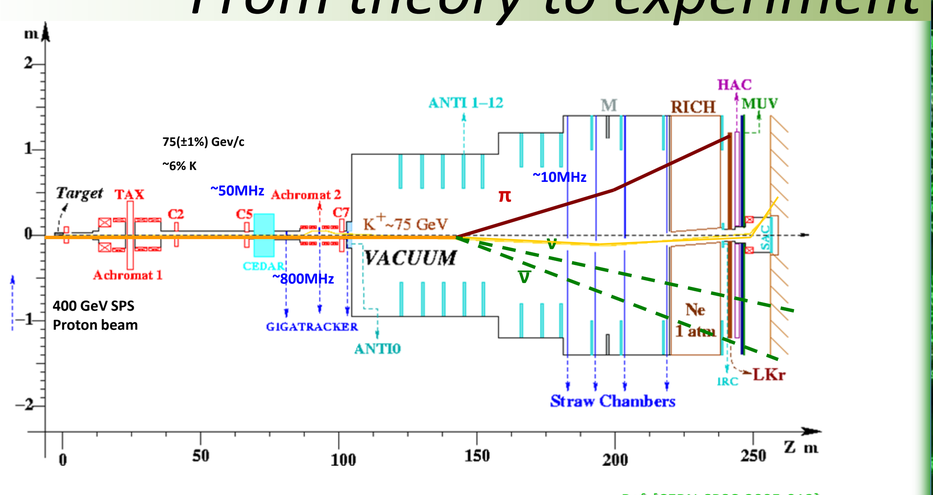

Ref. [CERN-SPSC-2005-013]

The NA62 trigger

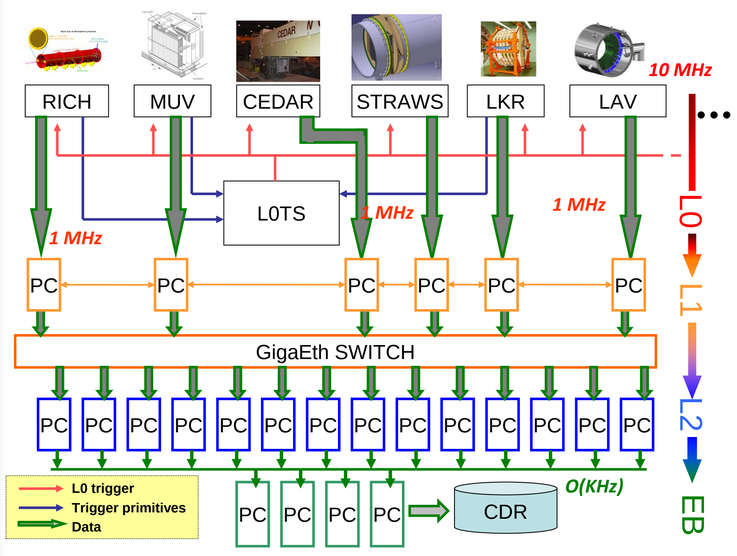

■ Multilevel trigger: high flexibility
L0 hardware, L1 and L2 software
■ The rate will be reduced from 10 MHz to few tens KHz
■ Integrated trigger and DAQ in a fully digital system
■ Fully monitored system: undetected acquisition losses $<10^{-8}$
■ The L0 central processor (L0TS)
L0 latency <1 ms (large buffers)

detector	Rate (MHz)
CEDAR	50
GTK	800
LAV (total)	9.5
STRAW (each)	8
RICH	8.6
LKR	10.5
MUV	9.2
SAC	1.5

The L0 bulding blocks

■ TELL1:generic DAQ board (EPFL Lausanne for LHCb exp.) used both for acquisition and generation of trigger primitives
■ Analog signals digitized by four 128-channels HPTDC boards: in total 512 channels on a single TELL1 board
■ Logic, primitives definitions, monitoring and buffering performed and controlled by 9 FPGAs for each board (5 on TELL1 and 4 onTDCBs) + 350 MB DDR RAM
■ TELL1 provides asynchronous GBE data transmission (4Gb/s) and receives synchronous trigger requests via TTC standard

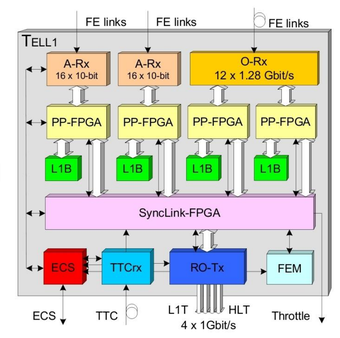

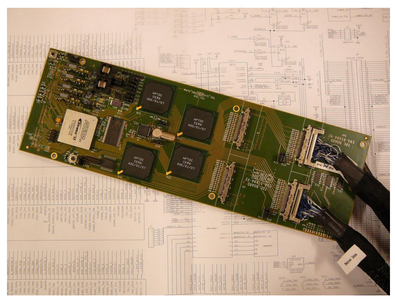

■ 128 TDC channels per board with 100 ps time resolution
■ QPLL on board to decrease the clock time jitter < 15ps
■ LVDS inputs
■ I2C and JTAG support
■ Altera Stratix II FPGA for data monitoring and preprocessing
■ 1MB SRAM
■ Low noise DC-DC converters and filters
■ Miniaturized connectors

RICH & Rings

■ RICH counter (Ne, 1atm) for particle identification downstream the decay region.
■ separation π–μ at $5x10^{-3}$ level (range 15-35 GeV/c)
■ Good time resolution: below 100 ps
■ ~2000 photomultipliers to identify Cherenkov rings
■ Partecipate to L0 trigger definition with multiplicity (by default)

■ Rings at trigger level → selective conditions
■ Ring's radius and center → velocity and angle
■ → coarse on-line measurement of particle momentum (with assumption on particle type)

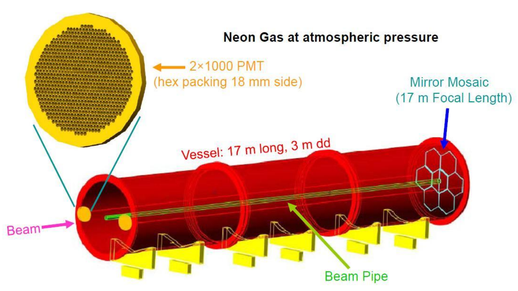

■ Ring finding at 10 MHz is a challenging task
■ On average one pion results in 20 hits
→ 1.6 GB/s (160 B/event)

Fast Ring finder

■ Simplest parallelized algorithm: one process for each possible center (each photomultiplier)
■ Without optimizations one ring found in <150 μs (stable timing results with negligible tails)
■ On track for use in <1ms latency on-line system

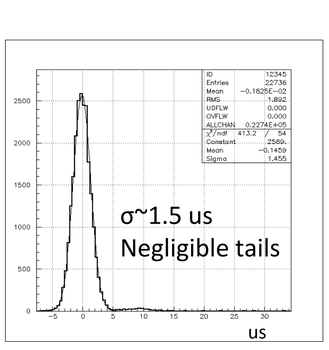

■ Goal: 100 μs per ring
■ best algorithm (under test):
Generalized Hough Transform

System used for tests: PC with
- CPU Intel i7-950 +12GB RAM,
- GPU: 2 x NVIDIA TESLA C1060
i.e. 2 x {240 cores + 4 GB VRAM}

GPU (Graphics Processor Unit)

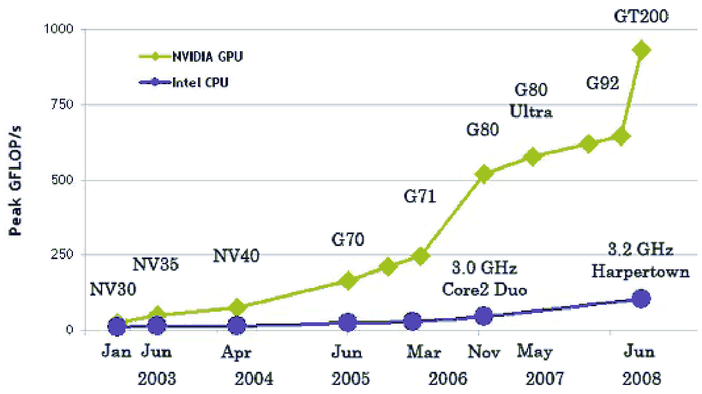

■ GPU = graphics processor in commercial Graphics Cards

■ GPUs power increasing faster than the CPUs due to:
(1) Different architecture:
more transistors for processing rather than caching, resource usage prediction and flow control
(2) Different scalability: the GPU power is related to the dimension and numbers of chips (more and more cores) and not to the chip speed (as in the CPU)

■ GPU operate simultaneous processing on thousands points:
large parallel computing power, large memory, huge bandwidths

■ Examples:

ATI-RADEON 4870:
10cores x16processors x 5sub-units
~1TFlops, ~100 GB/s bandwidth

NVIDIA-TESLA C1060:
30 cores x 8 processors + 4GB RAM
~1TFlops, ~100 GB/s bandwidth

GPUs	1 Tesla GPU
Single Precision Performance	933 Gigaflops
Double Precision Performance	78 Gigaflops
Memory	4 GB DDR3
Memory speed	800 MHz
Bandwidth	102 GB/s

■ Fast expanding field of scientific computation on GPUs (GPGPU)

■ This project: attempt to exploit GPUs in moderately large REAL-TIME environment

■ So far large latencies restricted this to high-performance computing; now O(100us) latencies easily obtained

■ Easy and cheap on a workstation
- a few TFlops → fast answers
- hundreds GB/s → high data rates

■ Fast links and real-time Linux for appropriate connection with the external world

■ The Graphics Cards sector is one of the most supported by IT industry (Digital imaging, video processing, Video Games,…): *can profit for free!*

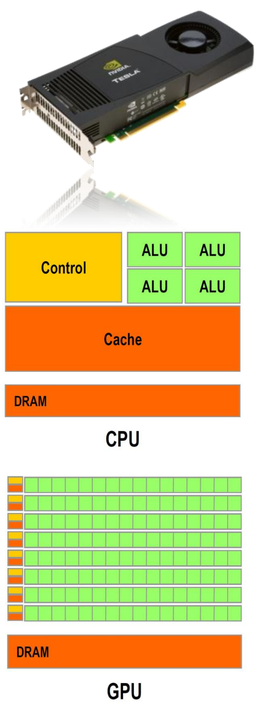

GPU integration in NA62 trigger

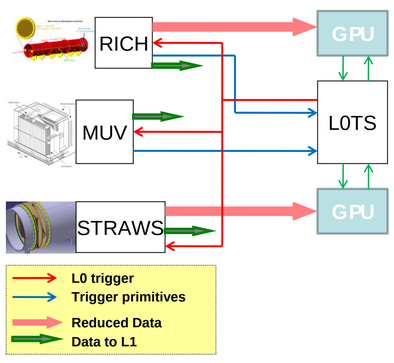

■ Simpler scheme for GPU cooperation with L0TS: links through PCI-e lanes or better SLI fast links

■ Other possibities in case of effective GPU performance not respecting the L0 requirements:
- L0.5: preselection by L0TS before sending to GPU
- L1: for decreasing the dimension of the L1 farm

■ Many of the detectors can partecipate to lowest trigger level thanks to the GPU approach (spectrometer, calorimeter, etc...)

■ Measurements with high resolution at L0 = great opportunity for efficient trigger in ultra high rate environment
■ Total trigger bandwidth for main physics channels (no zero-suppression) can be reduced by selective conditions early in trigger chain
■ TDAQ will not be a bottleneck in case of higher beam intensity for collecting more statistics
■ The scheme will give the unique possibility to collect additional interesting decays channels otherwise suppresed by the trigger for the main channels
■ Computing power pressure on the highr software trigger levels highly reduced

Gianmaria.collazuol@cern.ch, gianluca.lamanna@cern.ch, marco.sozzi@cern.ch