

Trigger/DAQ design: from test beam to medium size experiments

ISOTDAQ 2019
Royal Holloway, University of London
6 April 2019



few introductory notes ...

about the speaker ...

about the speaker ...

Roberto Ferrari
Istituto Nazionale di Fisica Nucleare

about the speaker ...

Roberto Ferrari

Istituto Nazionale di Fisica Nucleare

Oh my ! Yet another f...⁽¹⁾ Italian !⁽²⁾

about the speaker ...

Roberto Ferrari

Istituto Nazionale di Fisica Nucleare

Oh my ! Yet another f...⁽¹⁾ Italian !⁽²⁾

(1) fanatic ... friendly ... fantastic ... ?

(2) about 13.5 lectures (out of 30) covered by Italians

Students' homework

Debate one of the following hypothesis:

Students' homework

Debate one of the following hypothesis:

- 1) baseline: statistical fluctuation or new physics ?
(to be submitted to Nature)*

Students' homework

Debate one of the following hypothesis:

*1) baseline: statistical fluctuation or new physics ?
(to be submitted to Nature)*

*2) romantic: what about "Italians do it better" ?
(to be submitted to Vanity Fair)*

Students' homework

Debate one of the following hypothesis:

1) *baseline: statistical fluctuation or new physics ?*
(to be submitted to Nature)

2) *romantic: what about "Italians do it better" ?*
(to be submitted to Vanity Fair)

3) *last but not least (for the paranoid/complottist ones):*

Students' homework

Debate one of the following hypothesis:

*1) baseline: statistical fluctuation or new physics ?
(to be submitted to Nature)*

*2) romantic: what about "Italians do it better" ?
(to be submitted to Vanity Fair)*

*3) last but not least (for the paranoid/complottist ones):
what about the famous Mafia-Pizza-Spaghetti-TDAQ connection ?
(will go anonymous on the dark web)*

Seriously:

Seriously:

you can't afford such a **demanding** environment
without any specific training ...

Seriously:

you can't afford such a **demanding** environment
without any specific training ...

*... about the **Italians' way***

Seriously:

you can't afford such a **demanding** environment
without any specific training ...

*... about the **Italians' way***

luckily, on the web, there are plenty of survival kits

Example 1: basic course (mild concepts)

Example 1: basic course (mild concepts)

A Short Lexicon of Italian Gestures

For Italians, it comes naturally. But what do they mean when they talk with their hands?
Many things. Roll over the images to learn a few classic gestures. [Related Article »](#)



Perfect!



What in God's name
are you saying?



Nothing.



Someone talks too much.



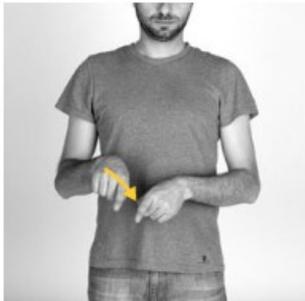
Get out of here.



Slow down or
keep calm.



I don't care.



Those two get along.



It wasn't me or
I don't know.



Don't worry,
I'll take care of it.



Why in God's name
did you/I do it?

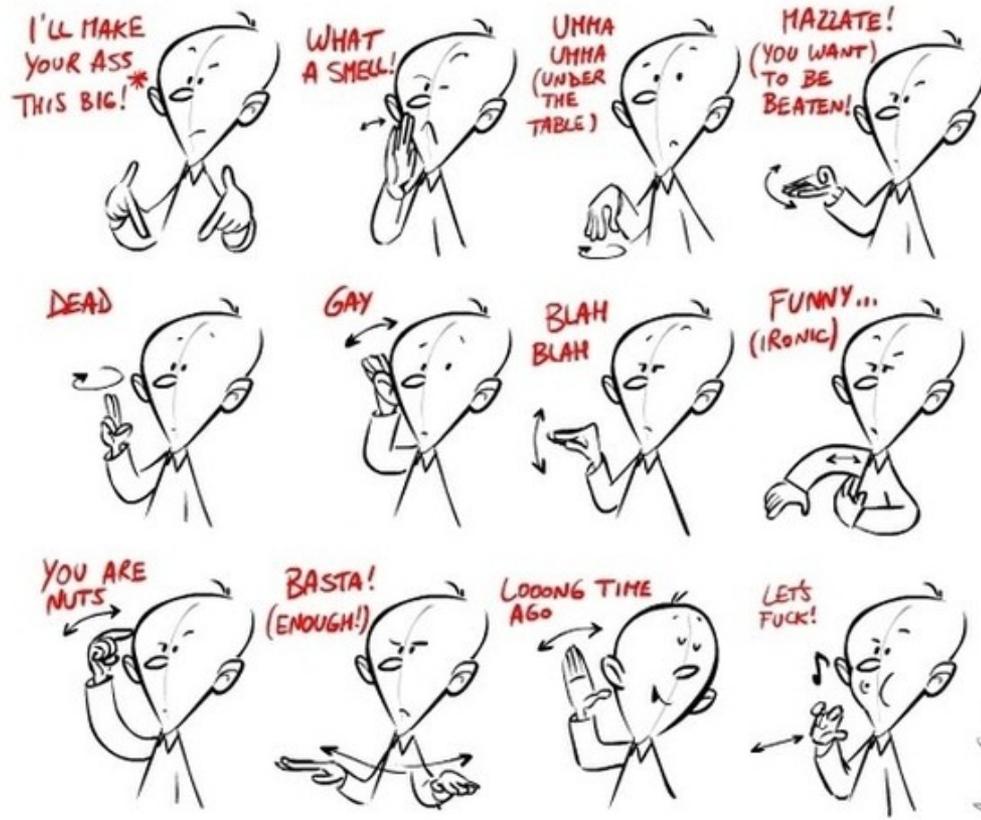


To be afraid.

Example 2: advanced course (includes sensitive concepts)

Example 2: advanced course (includes sensitive concepts)

- ITALIAN POPULAR GESTURES - CHAPTER II



* THE HANDS DISPLAY THE SIZE OF THE ANUS.

Please take care:

Please take care:

be careful while doing practice!

(especially for the advanced course)

more seriously ...

(1)

→ hope to give you something sensible ←

(2)

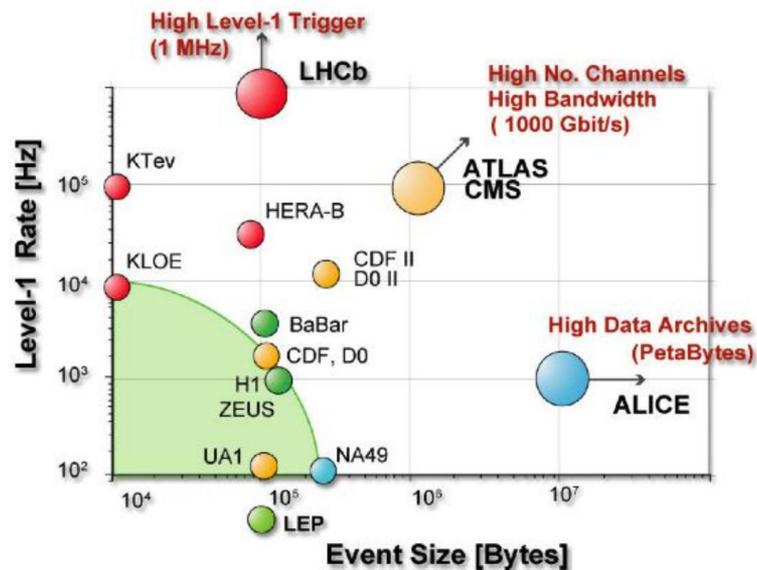
→ but, please, don't take anything at face value ←
just aiming at enlightening some critical issues

- not meant to be exhaustive (no way!) -

credit to Sergio Ballestrero most material from his talk at ISOTDAQ 2015



Trigger/DAQ design: from test beam to medium size experiments



Trying to move ...

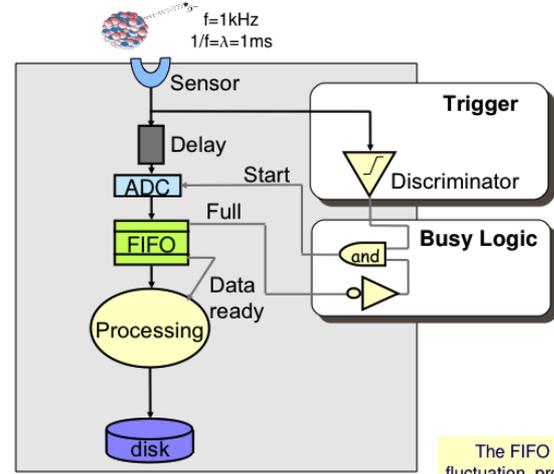
from here:



to here:



Basic DAQ: De-randomization



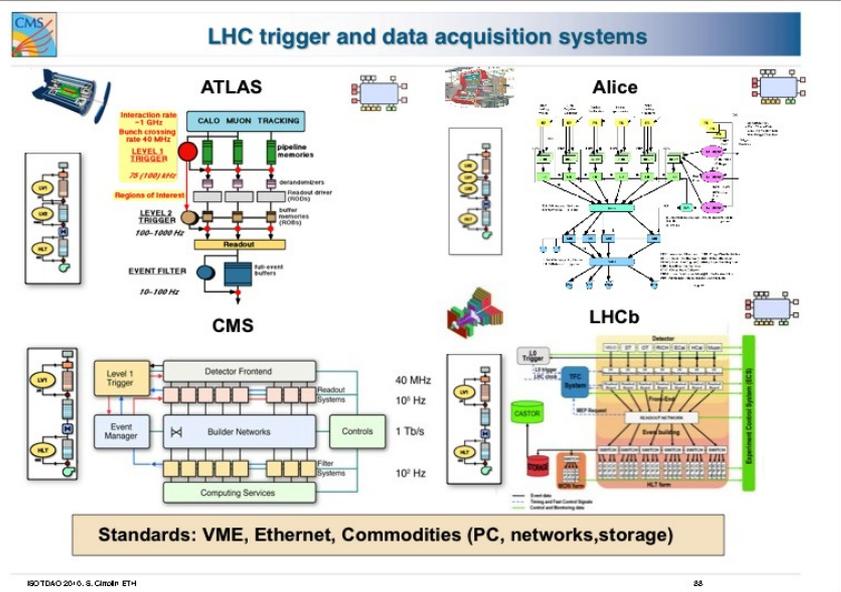
→ First-In First-Out

- Buffer area organized as a queue
- Depth: number of cells
- Implemented in HW and SW

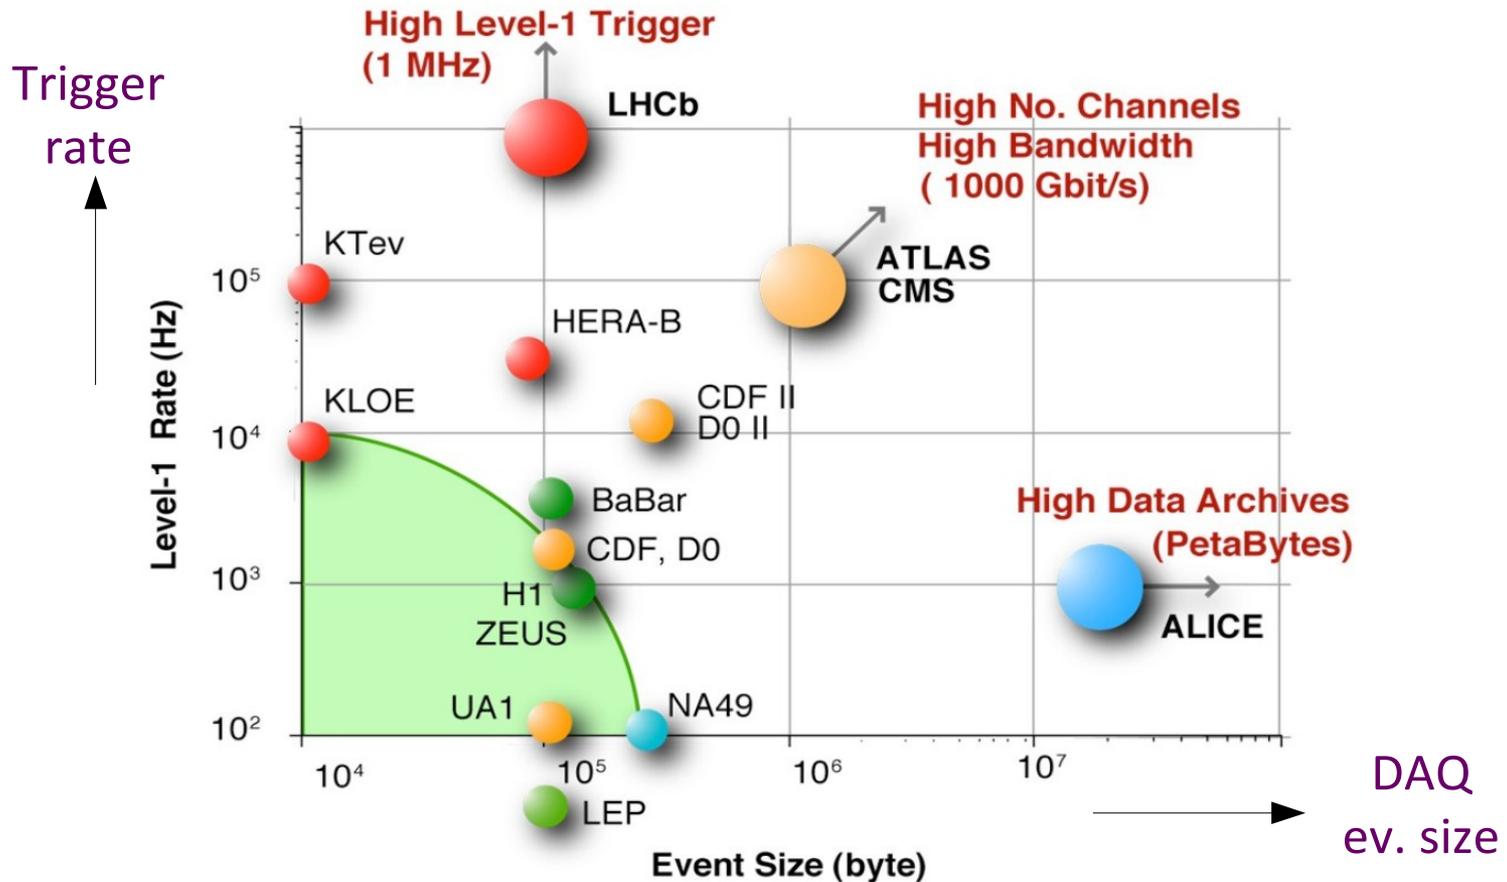


→ FIFO introduces an additional latency on the data path

The FIFO absorbs and smooths the input fluctuation, providing a ~steady (De-randomized) output rate

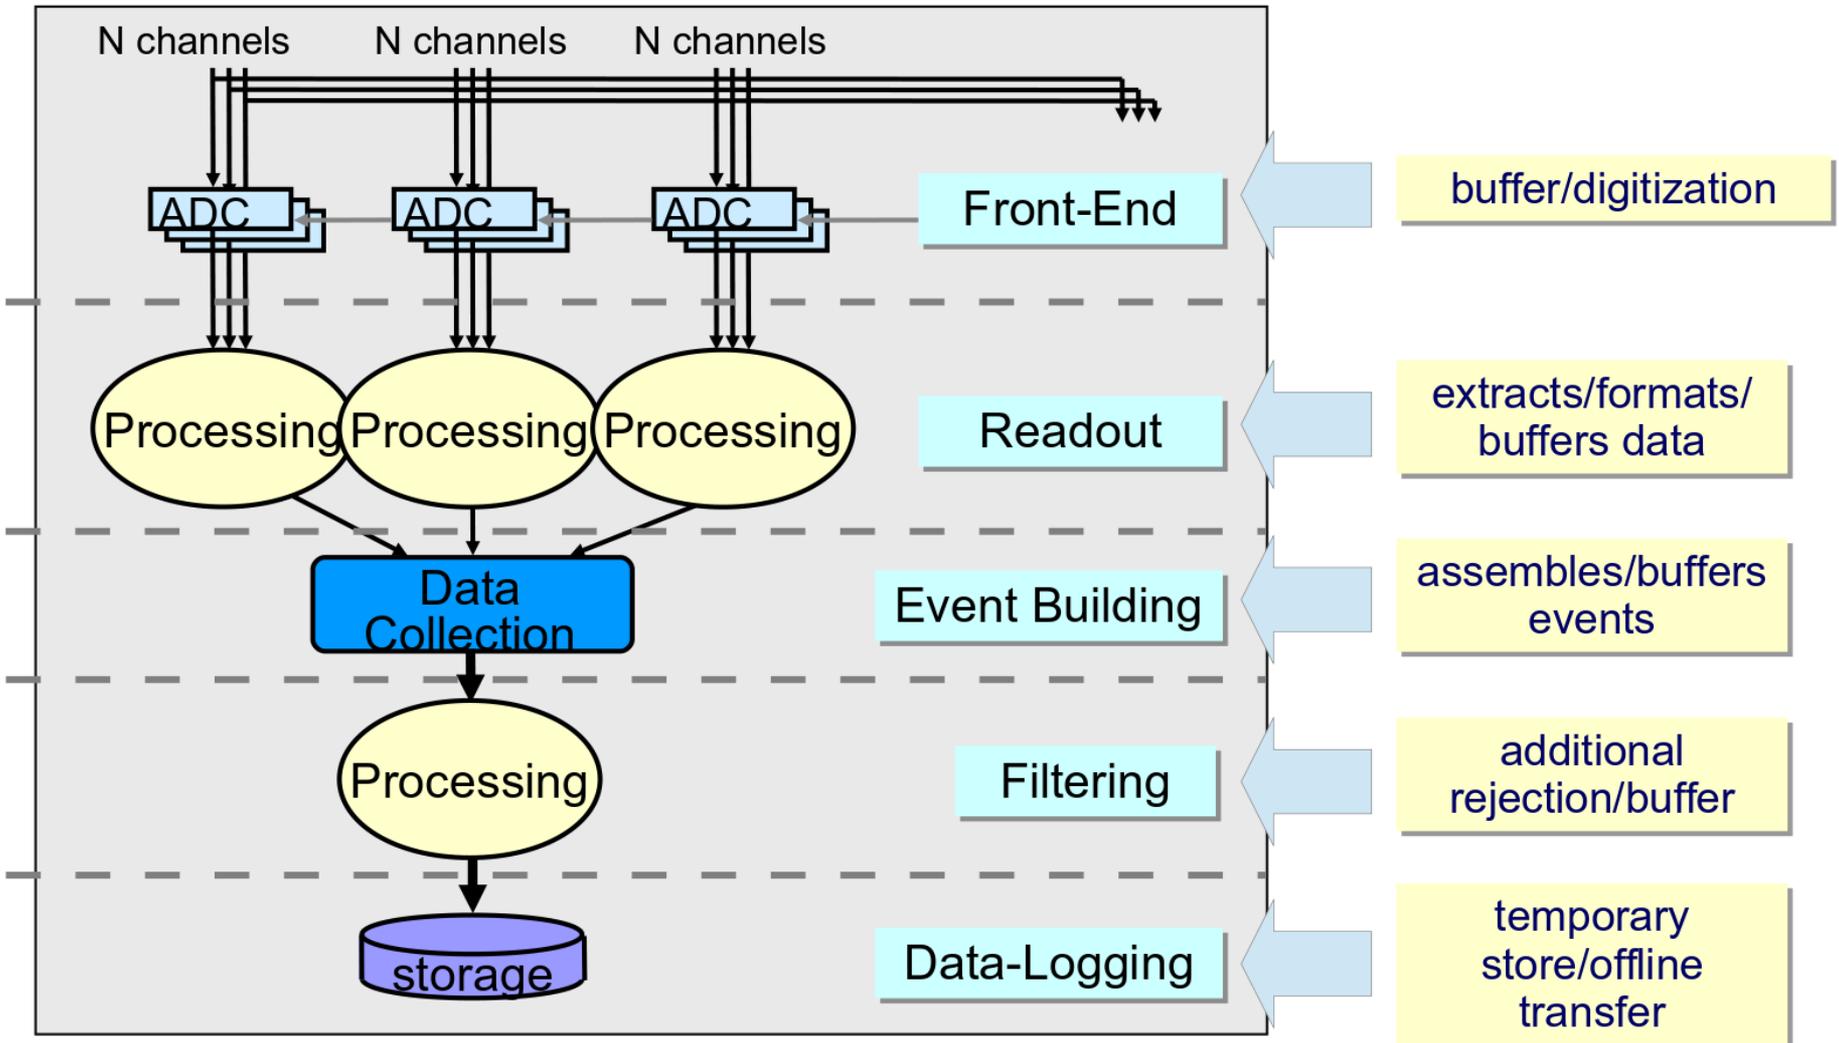


Trigger & DAQ in HEP



different issues → different solutions
no magic, unique solution for all cases

medium/large DAQ: constituents



aiming to cover the gap in 5 steps ...

- Step 1: Increasing the rate
- Step 2: Increasing the sensors
- Step 3: Multiple Front-Ends
- Step 4: Multi-level Trigger
- Step 5: Data-Flow control

back to square one

Do we really need a trigger ?

back to square one

Do we really need a trigger ?

not obvious ... triggerless DAQ systems do exist

back to square one

Do we really need a trigger ?

not obvious ... triggerless DAQ systems do exist

even in high-energy physics

back to square one

Do we really need a trigger ?

not obvious ... triggerless DAQ systems do exist

even in high-energy physics, e.g.:

a) LHCb upgrade 40 MHz readout

→ *briefly described in Francesca's final talk*

b) DUNE LAr TPC 2 MHz readout

→ *Giovanna's talk on April 11*

but triggering may be crucial ...

https://en.wikipedia.org/wiki/Coincidence_circuit :

Walther Bothe (1924-1929):

offline → online coincidence (logic **AND**) of 2 signals

Bruno Rossi (Nature, 1930):

"Method of Registering Multiple Simultaneous Impulses of
Several Geiger Counters"

→ online coincidence of 3 signals (scalable)!

first modern trigger

https://en.wikipedia.org/wiki/Coincidence_circuit :

“Rossi coincidence circuit was rapidly adopted by experimenters around the world. It was the first practical AND circuit, precursor of the AND logic circuits of electronic computers”

Geiger-Muller
counters

Rossi's circuit: coincidence of
signals of 3 Geiger-Muller
counters

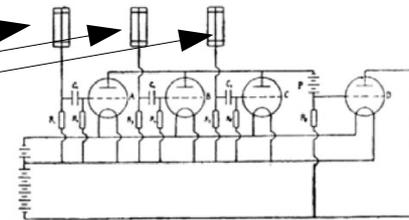
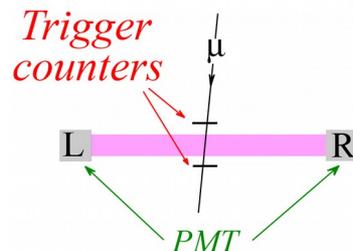


Fig. 17 – Il circuito di Rossi per rivelare coincidenze di raggi cosmici che arrivano sui contatori Geiger (i rettangoli in alto dello schema)¹⁹.

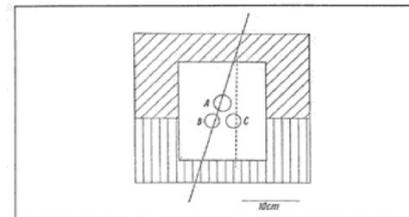


Fig. 18 – L'uso del circuito di Rossi per rivelare una coincidenza tripla che, nella disposizione in figura dei tre contatori, mostra la produzione di una radiazione secondaria (linea tratteggiata) da parte della radiazione primaria (linea continua)²⁰.

simplest case: 2-signal
coincidence

but even simple trigger systems ...

Gokhan's talk:

$$N1 = s1 \cdot s2 \cdot \overline{s3}$$

Veto (anti-coincidence)

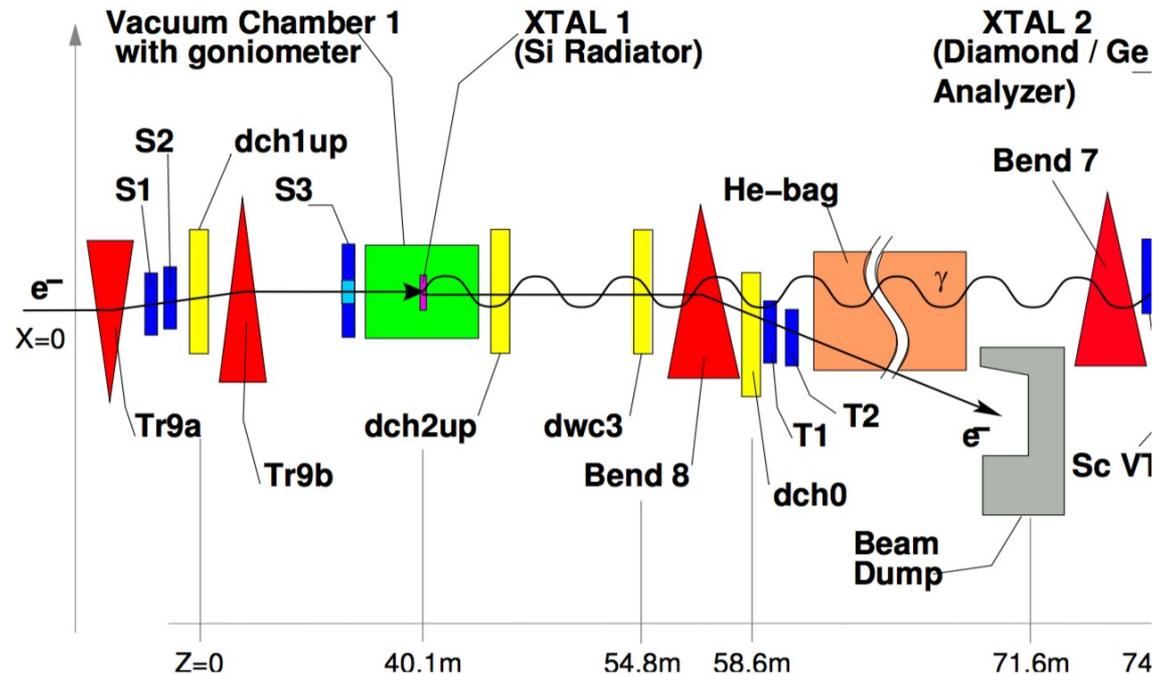
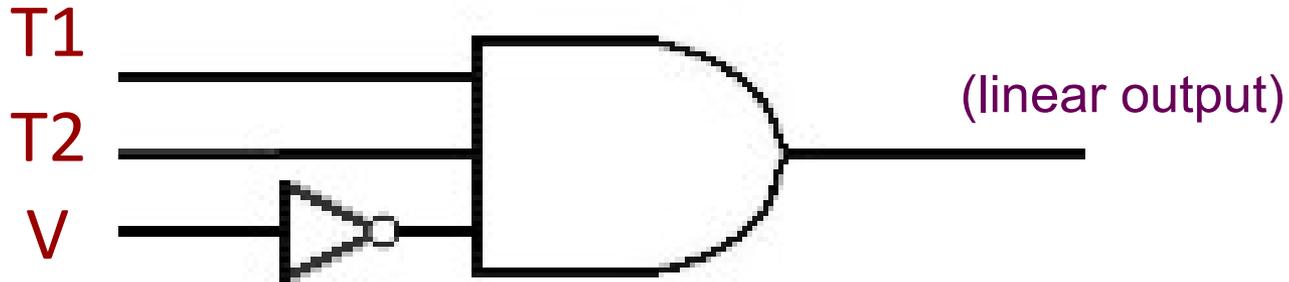


Fig. 1. Setup of the Na59 Experiment

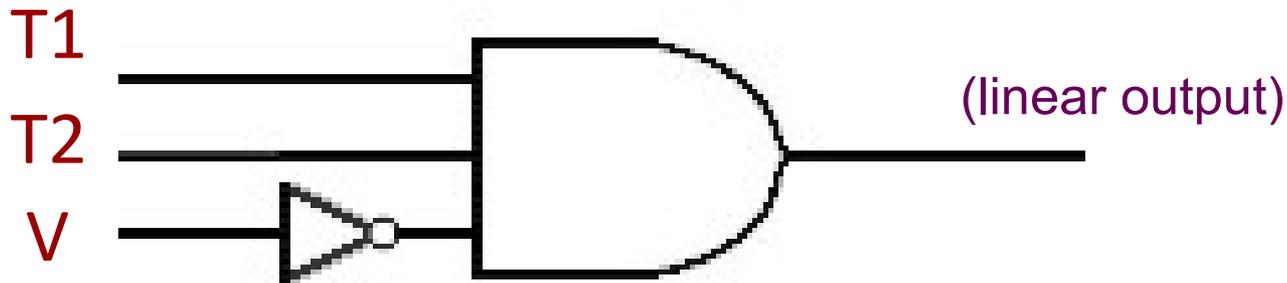
... pose issues !

... pose issues !



(anti-)coincidence with veto
→ simple, clear !

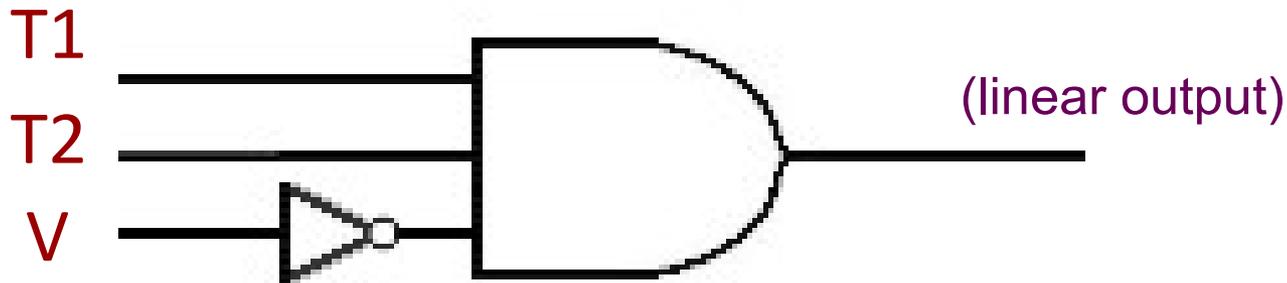
... pose issues !



(anti-)coincidence with veto
→ simple, clear !

really doing what you think/need ?

(anti-)coincidence with veto



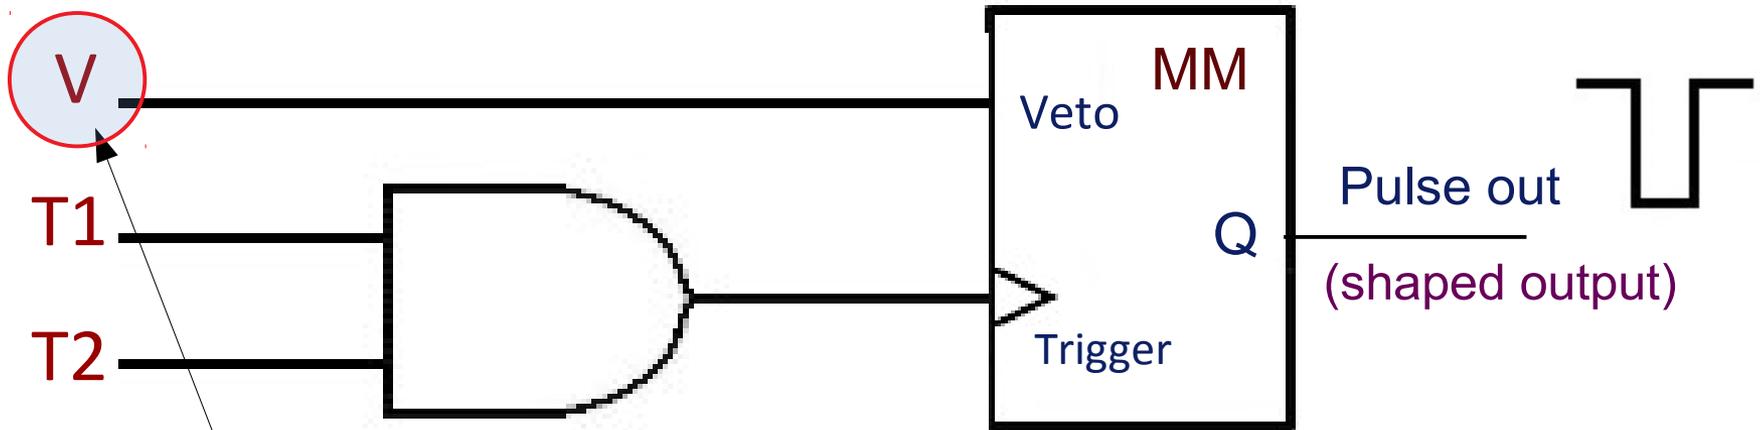
flawed !

output signal may:

- a) jitter
- b) fluctuate in duration
(or both)

because of relative timing of T1, T2, V

(anti-)coincidence with veto



can be a busy signal

MM = Monostable Multivibrator
= One-Shot Pulse Generator

much better !

first lesson(s)

trigger signal:

1) should be formed

→ pulse with predefined duration

2) veto/busy should block pulse generation

step one: increase rate

Many issues:

- trigger latency
- readout latency
- throughput
- rate fluctuations (trigger bursts)
- throughput fluctuations
(correlated noise, ...)

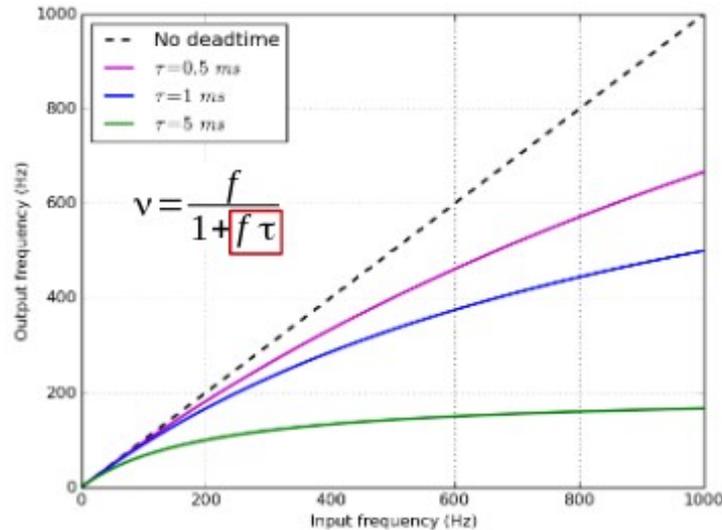
step one: increase rate

Many issues:

- trigger latency
- readout latency
- throughput
- rate fluctuations (trigger bursts)
- throughput fluctuations
(correlated noise, ...)
- dead-time

dead time (from Andrea's introduction)

Deadtime and efficiency

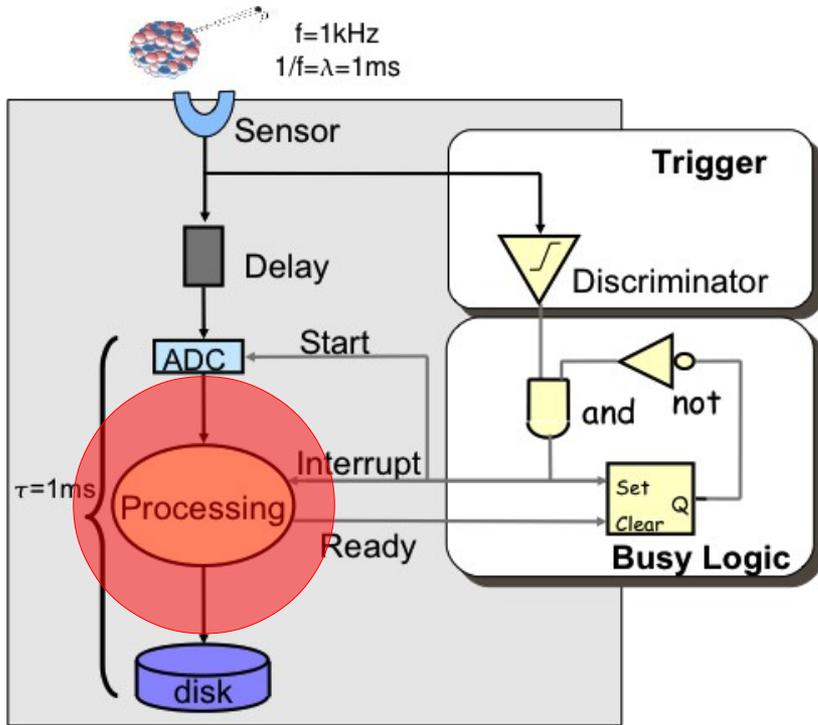


- In order to obtain $\epsilon \sim 100\%$ (i.e.: $v \sim f$) $\rightarrow f\tau \ll 1 \rightarrow \tau \ll \lambda$
 - E.g.: $\epsilon \sim 99\%$ for $f = 1\text{ kHz}$ $\rightarrow \tau < 0.01\text{ ms}$ $\rightarrow 1/\tau > 100\text{ kHz}$
 - To cope with the input signal fluctuations, we have to **over-design** our DAQ system by **a factor 100**
- How can we mitigate this effect?



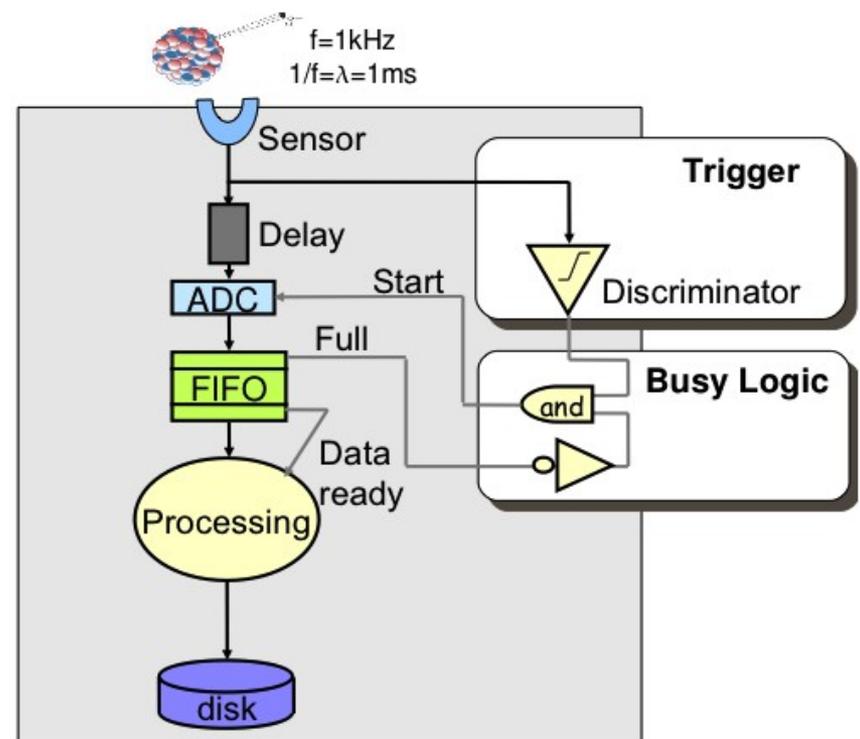
dead time → de-randomise

- Processing → bottleneck ?



$(f \cdot \tau) \sim 1 \rightarrow$ dead time $\sim 50\%$

- Buffering → decouple problems



What the impact ?

$(f \cdot \tau) \sim 1 \rightarrow$ dead time ?

buffering solve all problems ?

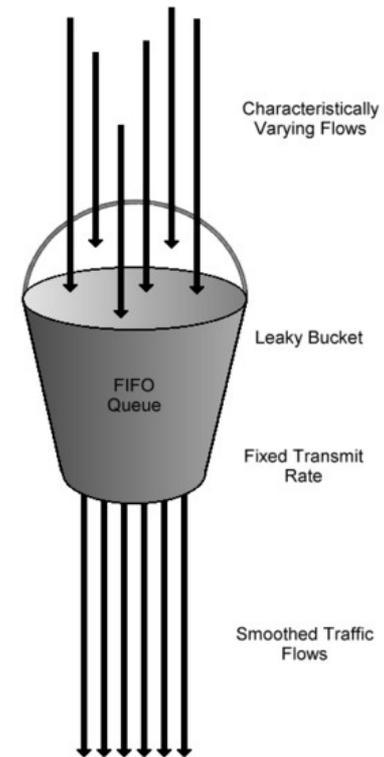
- FIFO

1) filling at very variable input flow

2) emptying at smoothed output flow

→ the Leaky-Bucket problem

Q: how often may overflow ?



off-topic: some very basic queueing theory

- N-event buffer ... single queue size N:

P_k : % time with k events in $\Rightarrow P_N = \text{no space available} \rightarrow \text{dead time}$

$$\sum P_k = 1 \quad [k=0..N]$$

$$\text{rate}(j \rightarrow j+1) = f \cdot P_j \quad (\text{fill at rate } f)$$

$$\text{rate}(j+1 \rightarrow j) = v \cdot P_{j+1} \quad (\text{empty at rate } v \gtrsim f)$$

$$\text{steady state: } v \cdot P_{j+1} = f \cdot P_j \Rightarrow P_{j+1} = \rho \cdot P_j = \rho^{j+1} \cdot P_0 \quad \text{where } \rho = (f/v) \lesssim 1$$

$$\text{for } \rho \sim 1 \Rightarrow P_j \sim P_{j+1} \Rightarrow \sum P_k \sim (N+1) \cdot P_0 = 1 \Rightarrow P_0 \sim P_N \sim 1/(N+1)$$

$$\Rightarrow \text{dead time} \sim 1/(N+1)$$

$$\text{want d.t.} \lesssim 1\% \Rightarrow N \gtrsim 100$$

off-topic: some very basic queueing theory

- N-event buffer ... single queue size N:

P_k : % time with k events in $\Rightarrow P_N = \text{no space available} \rightarrow \text{dead time}$

$$\sum P_k = 1 \quad [k=0..N]$$

$$\text{rate}(j \rightarrow j+1) = f \cdot P_j \quad (\text{fill at rate } f)$$

$$\text{rate}(j+1 \rightarrow j) = v \cdot P_{j+1} \quad (\text{empty } \approx f)$$

steady state: $v \cdot P_{j+1} = f \cdot P_j \Rightarrow P_{j+1} = \rho \cdot P_j = \rho^{j+1} \cdot P_0$ where $\rho = (f/v) \lesssim 1$

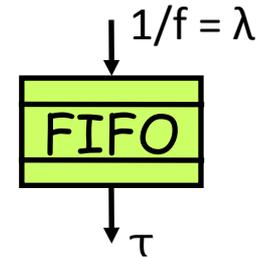
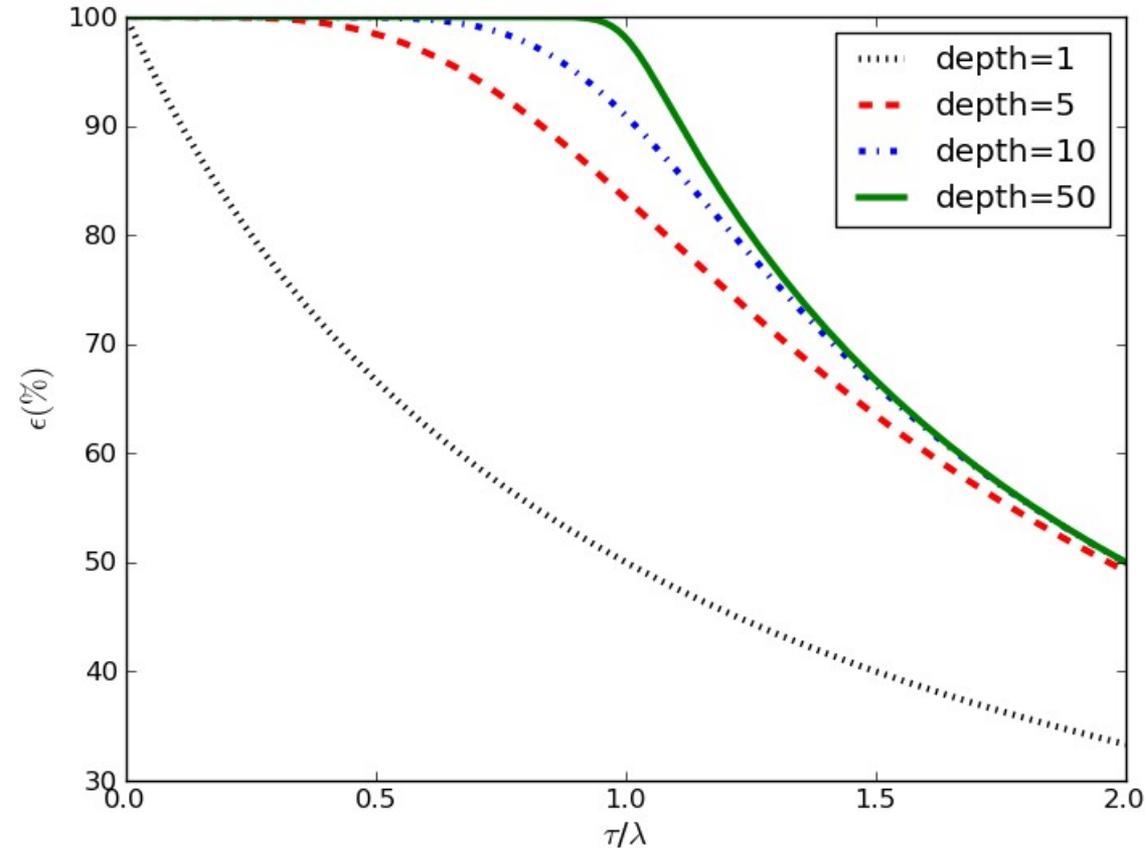
$$P_j \sim P_{j+1} \Rightarrow \sum P_k \sim (N+1) \cdot P_0 = 1 \Rightarrow P_0 \sim P_N \sim 1/(N+1)$$

$\Rightarrow \text{dead time} \sim 1/(N+1)$

want d.t. $\lesssim 1\% \Rightarrow N \gtrsim 100$

Take care: analytic calculation possible for pretty simple systems only

de-randomisation



- We can now attain a DAQ efficiency $\sim 100\%$ with:
 - $\tau \sim 1/f$
 - “moderate” buffer size
- One more degree of freedom to play with
- Often managed by trigger system itself (“complex dead time”)

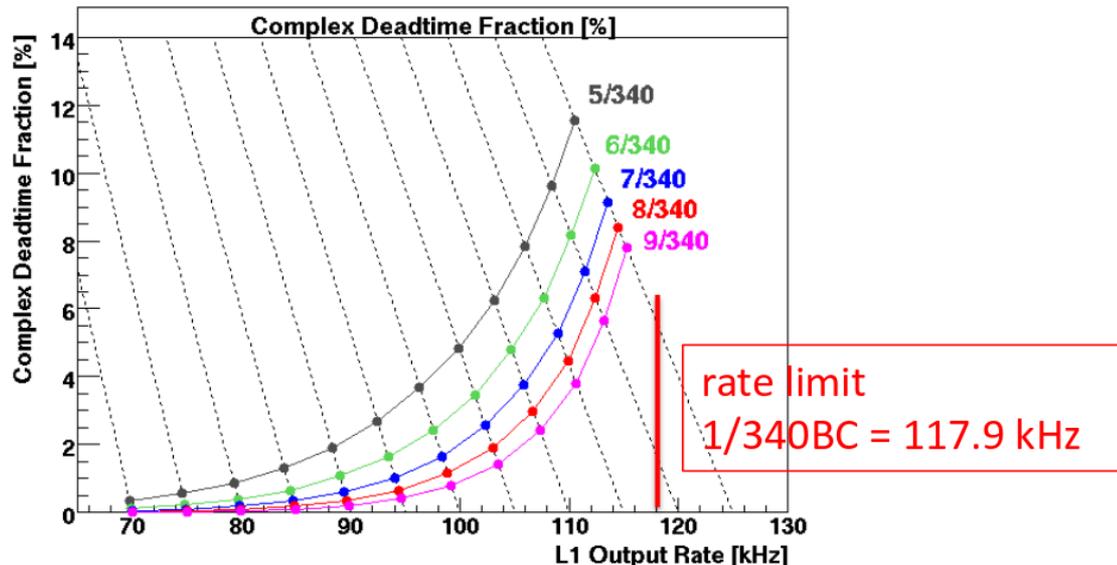
complex dead time

1) simple dead time: avoid overlapping (conflicting) readout window

2) complex dead time: avoid overflow in front-end buffers (protection against trigger bursts)

e.g. ATLAS: simple leaky-bucket algorithms with 2 parameters:

max X triggers (X = FIFO depth) in any (sliding) time window = (X*readout time)



FIFO

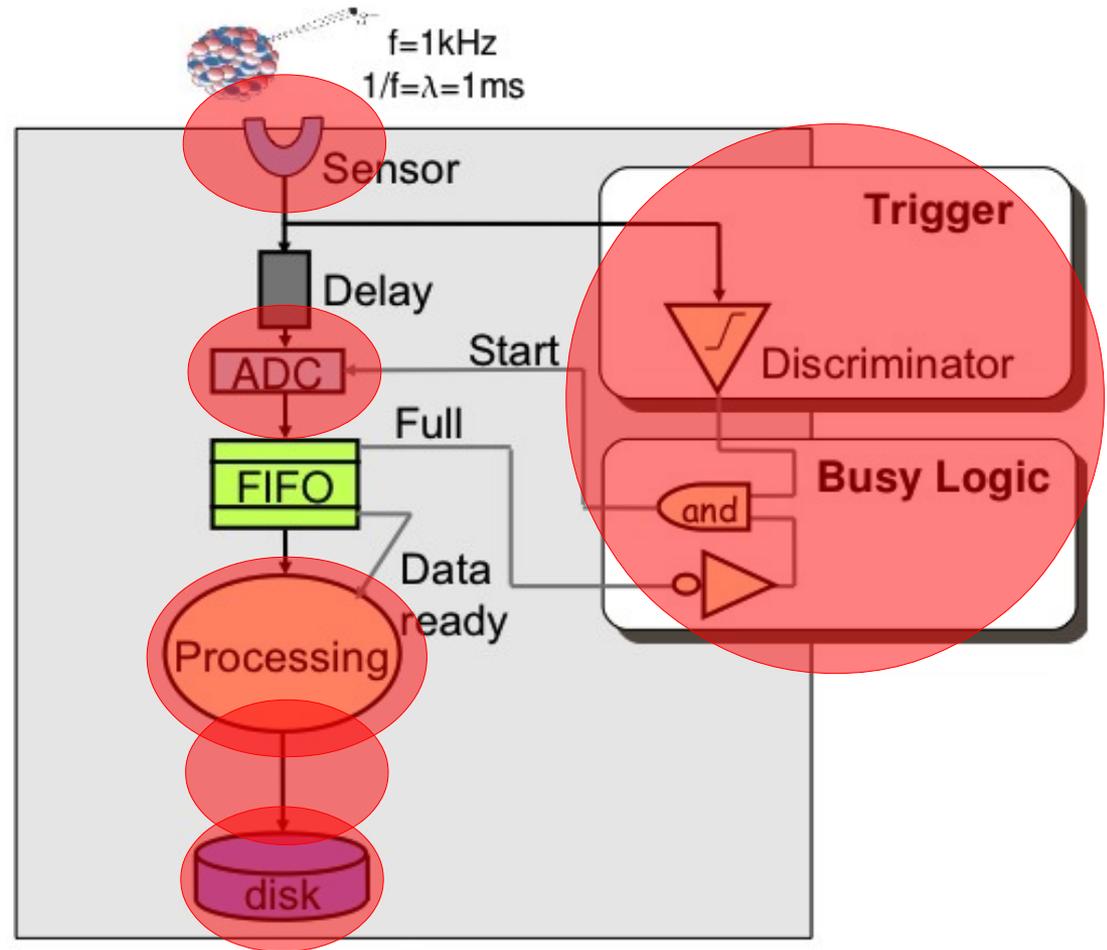
First-In First-Out memory:

- 1) independent read/write (sequential) access
- 2) may be hardware or over RAM

better with Dual-Port RAM

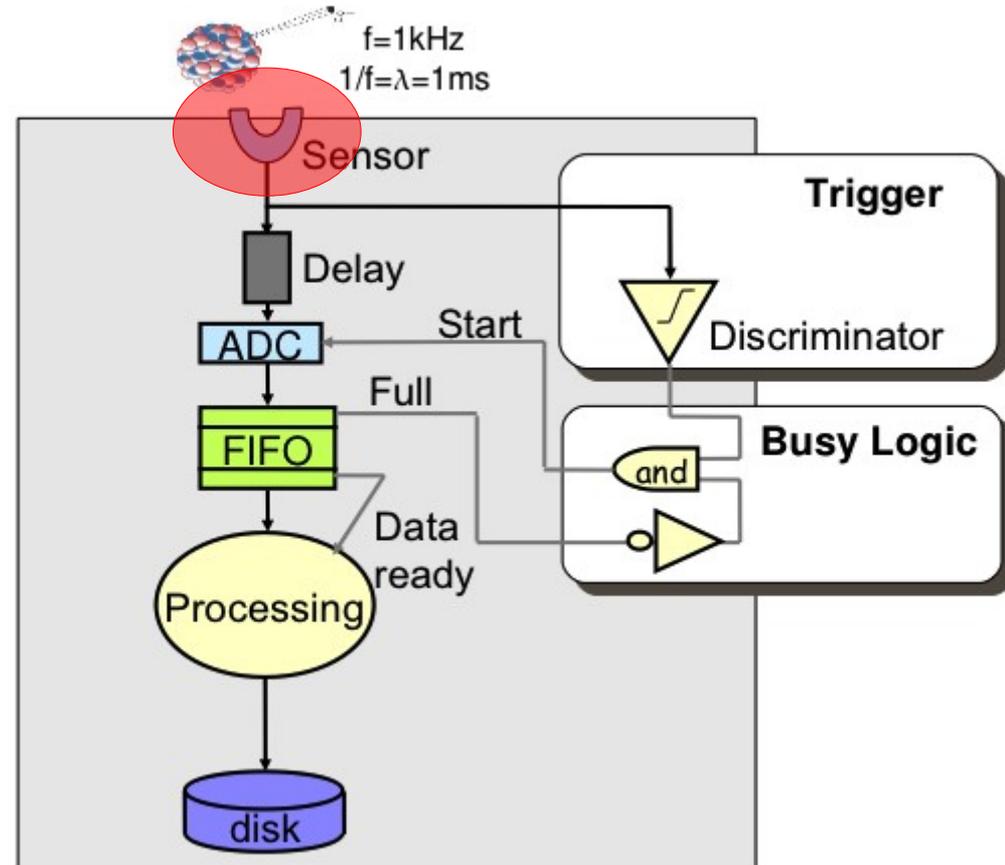
game over ?

many other possible limits even in a simple DAQ



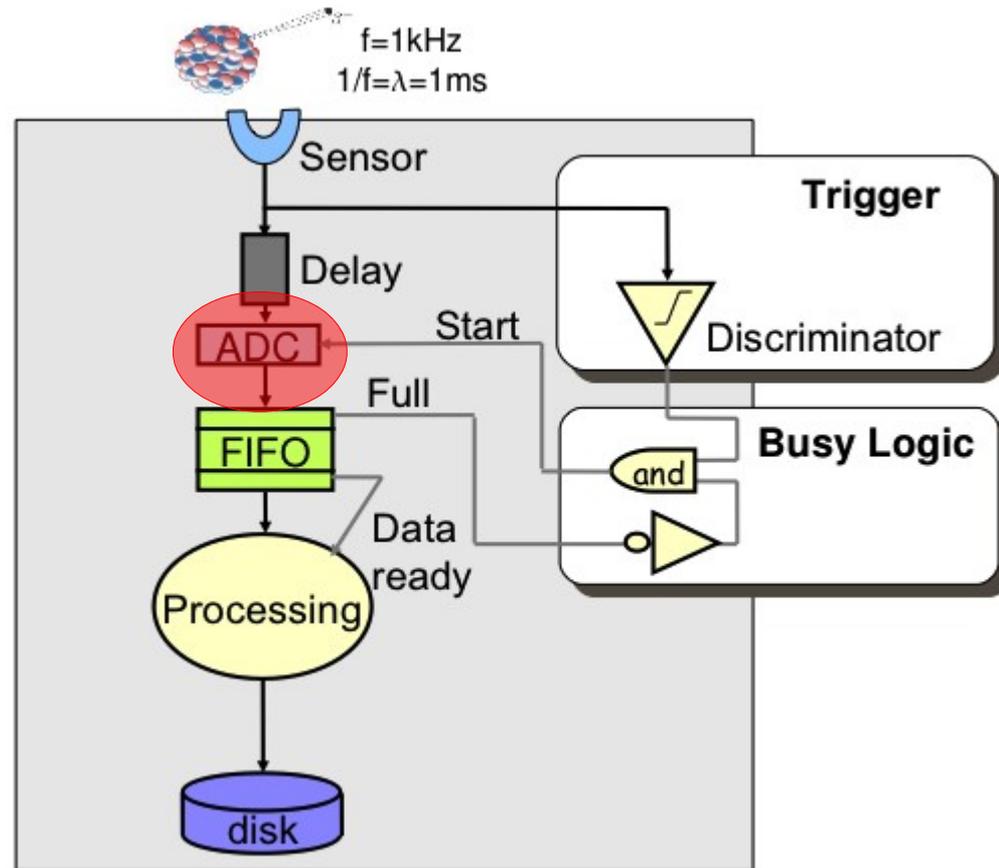
→ sensor

- Sensors limited by physical processes such as:
 - drift times in gases
 - charge collection in Si
- (possibly) choose fast processes
- analog FE imposes limits as well
- split sensors, each gets less rate:
“increase granularity”



→ ADC

- A/D conversion also limited
- Fast ADC
 - # of bits (resolution)
 - power consumption
- Alternatives:
 - analog buffers
 - (e.g. switched capacitor arrays)
- You may need integration (or sampling) over quite some time

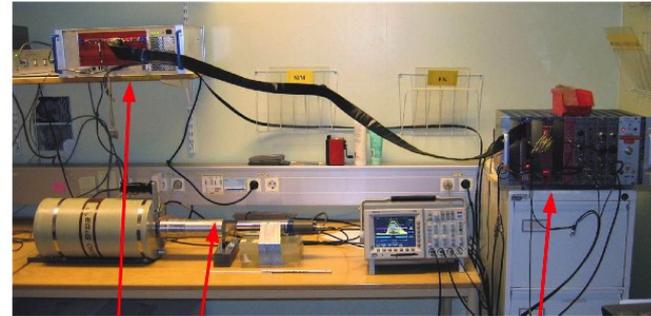


an example

- HPGe + NaI Scintillator
High res spectroscopy and beta+ decay identification
- minimal trigger with busy logic
- Peak ADC with buffering, zero suppression
- VME SBC with local storage
- 3×12 bits data size
+ 32 bit timestamp
- Root for monitor & storage
- Rate limit ~14 kHz
 - HPGe signal shaping for charge collection
 - PADC conversion time



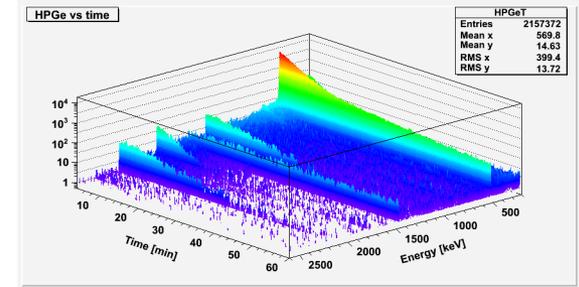
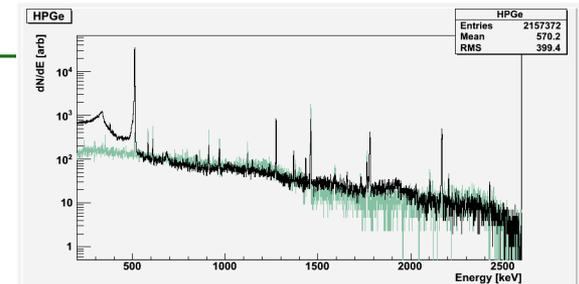
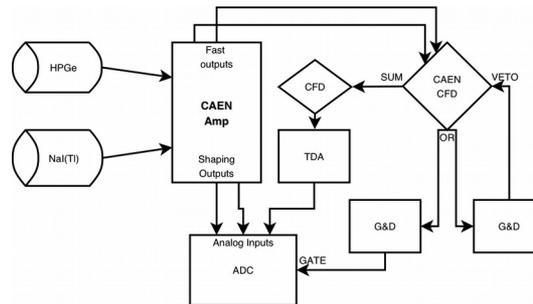
Ge crystal for isotope identification



Crystal HPGe

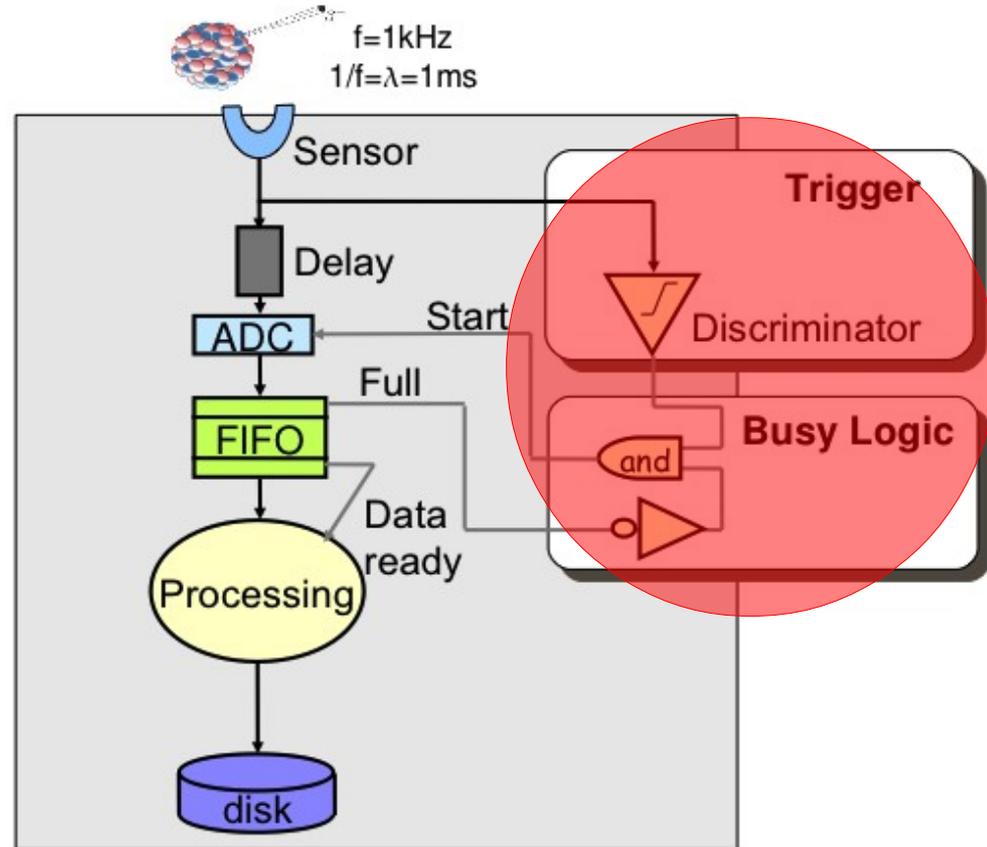
Trigger & front-end

Readout (ADC)



→ trigger latency

- simple trigger: ~fast
- complex trigger logic: not obvious [even when all in hw]
- some trigger detectors may be far away / slow → latency
- trigger signal is one: all information at a single point
 - in one step:
too many cables
 - in many steps:
delays



→ discrete modules: ~ 5-10 ns delay → tot. latency \geq 20-30 ns ←

DREAM (2006→): a testbeam case

R&D on dual-readout calorimetry, setup:

- Crystals
- Scintillating/cherenkov fibers in lead/copper matrices
- Scintillator arrays as shower leakage counters
- Trigger/veto/muon counters
- Precision chamber hodoscope

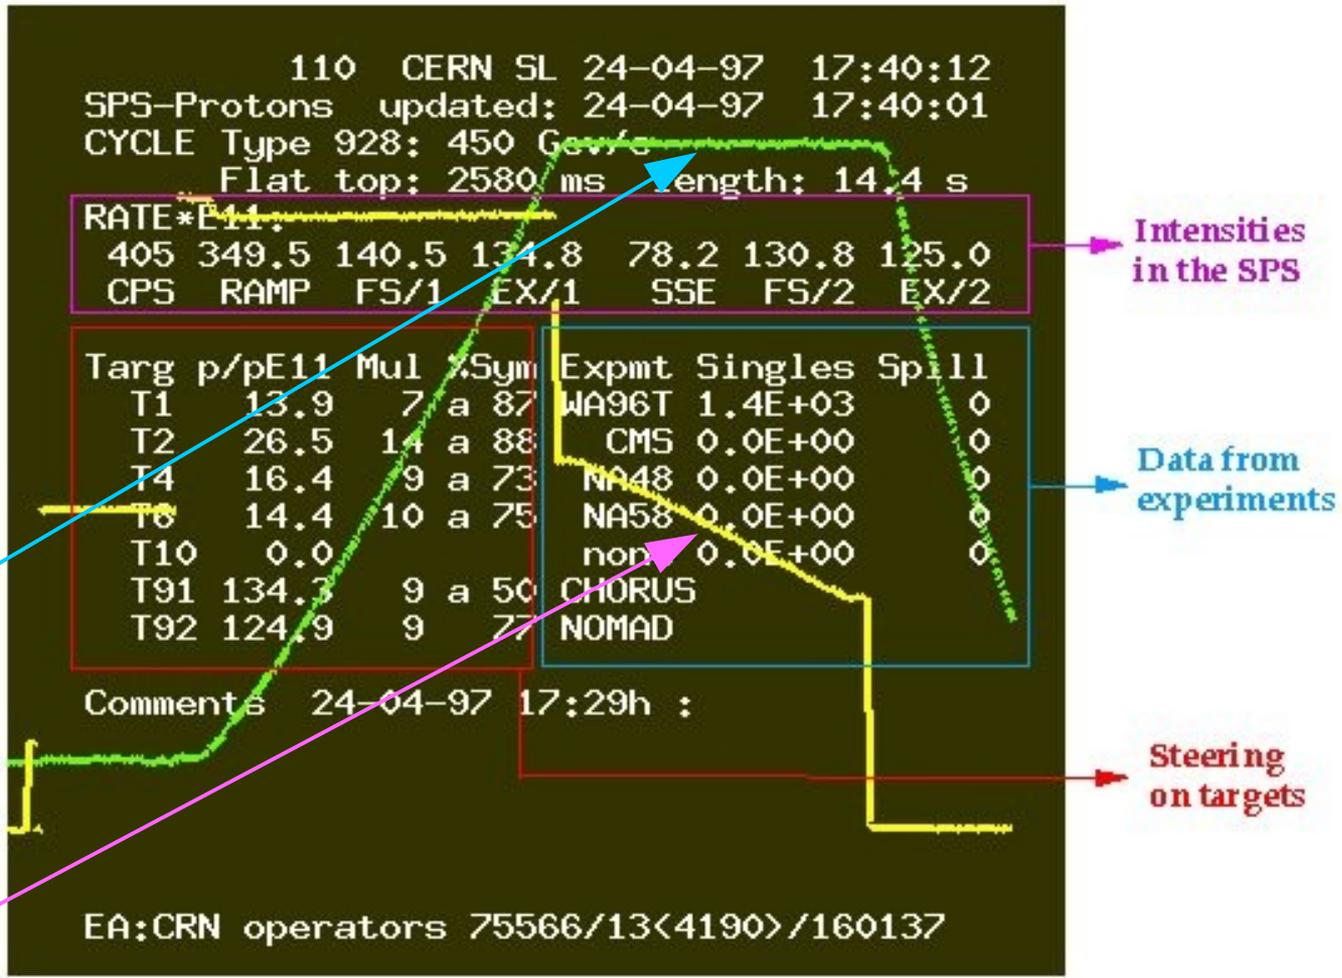
... always evolving

Acquiring: waveforms, total charge, time information

DREAM (2006→): a testbeam case

a possible
SPS cycle

duty cycle:
~ 2 s / 14.4 s
(flat top)



$$\text{Trigger} = (\overline{V} \times T_1 \times T_2 \mid \text{ped}) \rightarrow \text{easy!}$$

readout system

1 PC → readout of 2 VME crates (via CAEN optical interfaces)

1 PC → storage

6 x 32 ch xDC.s (x = Q, T : CAEN V792, V862, V775)

1 x 34 ch (CAEN V1742) 5 Gs/s Digitizer
(single event: $\sim 34 \times 1024 \times 12$ bit)

1 x 4 ch Tektronix TDS7254B 20 Gs/s oscilloscope

... few VME I/O & discriminator boards

... all in the control room

dataflow

1) Pull mode → FE electronics waiting for PC readout
(self-blocking trigger, re-enabled after readout)

2) Block data transfer → DMA (Direct Memory Access)
data moved by specialised hw (not by CPU)

[Push mode → FE electronics sending data as soon as
available]

DAQ

DAQ logic spill-driven (no “real time”, SLC desktops)

in-spill (slow extraction)

poll trigger signal ... if trigger present:

- a) (block) read all VME boards
- b) format & store on large buffers (FIFO over RAM)
- c) re-enable trigger

out-of-spill

- a) read scope (in case) → event size fixed at run start
- b.1) flush buffers to disk (beam and pedestal files) over network
- b.2) monitor data (produce root files)

rate $\sim O(1 \text{ kHz})$

trigger system

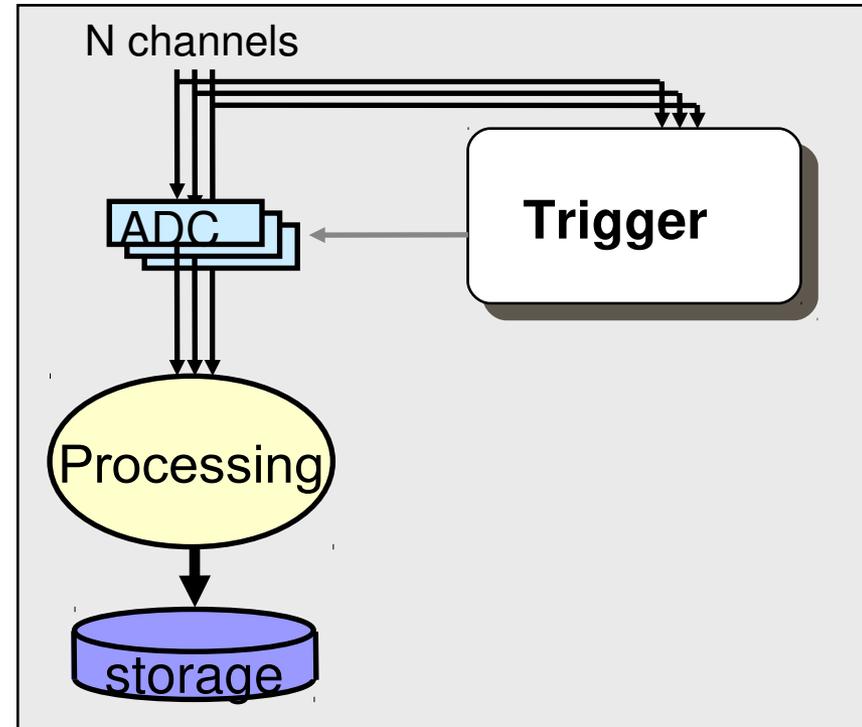
- a) crystals w/ fast PMT.s
- b) no analog buffering

→ low-latency trigger

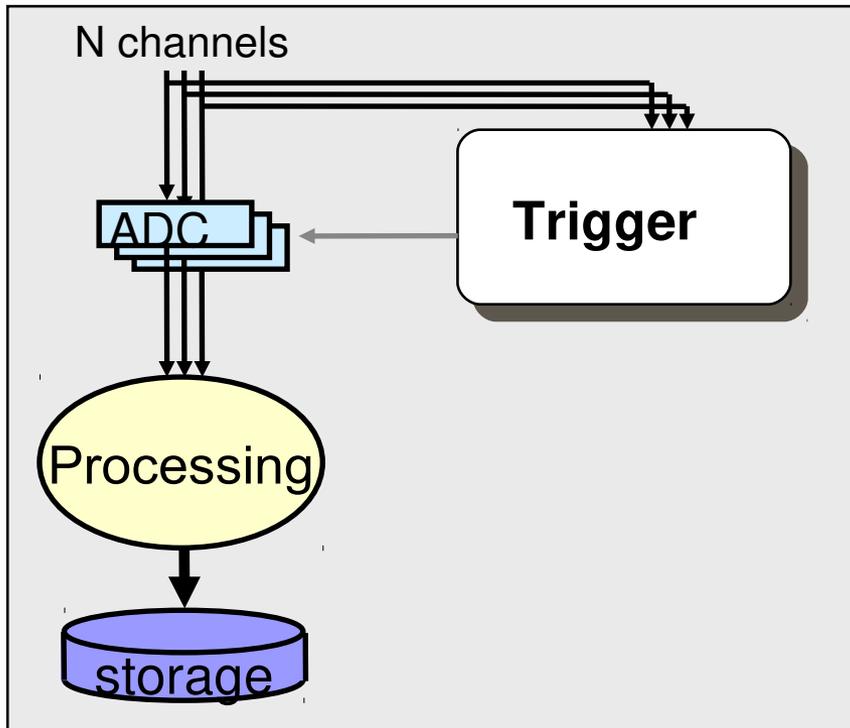
first discrete, then FPGA
(Xilinx Spartan 3AN evaluation board)

step two: increase # of sensors

- More granularity at the physical level
- Multiple channels (usually with FIFOs)
- Single, all-HW trigger
- Single processing unit
- Single I/O

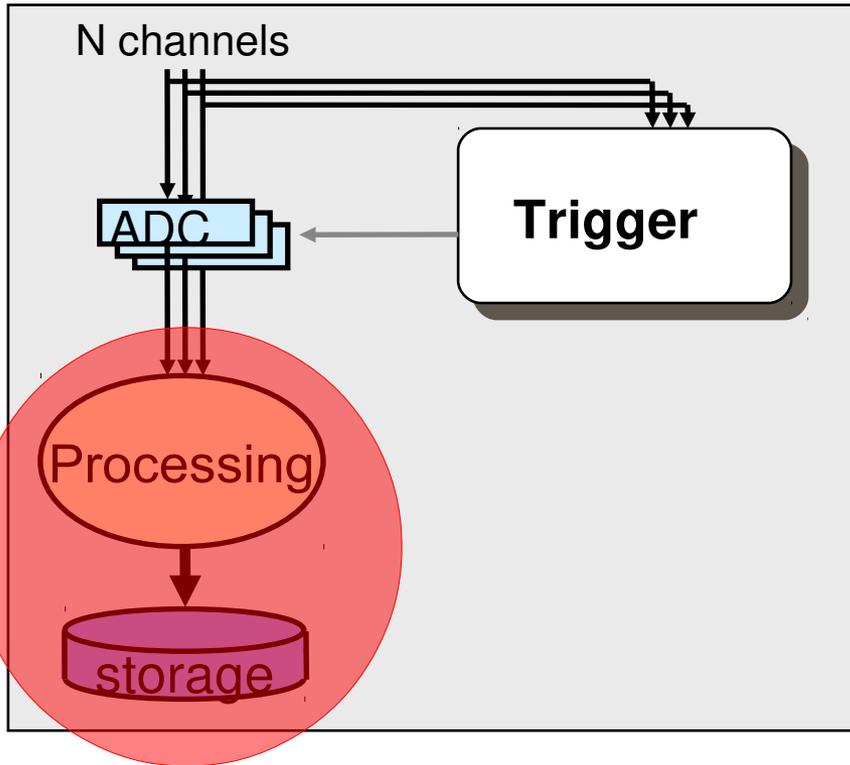


multi-channel, single-PU system



- common architecture in test beams and small experiments
- often rate limited by (interesting) physics itself, not TDAQ system
- or by the sensors

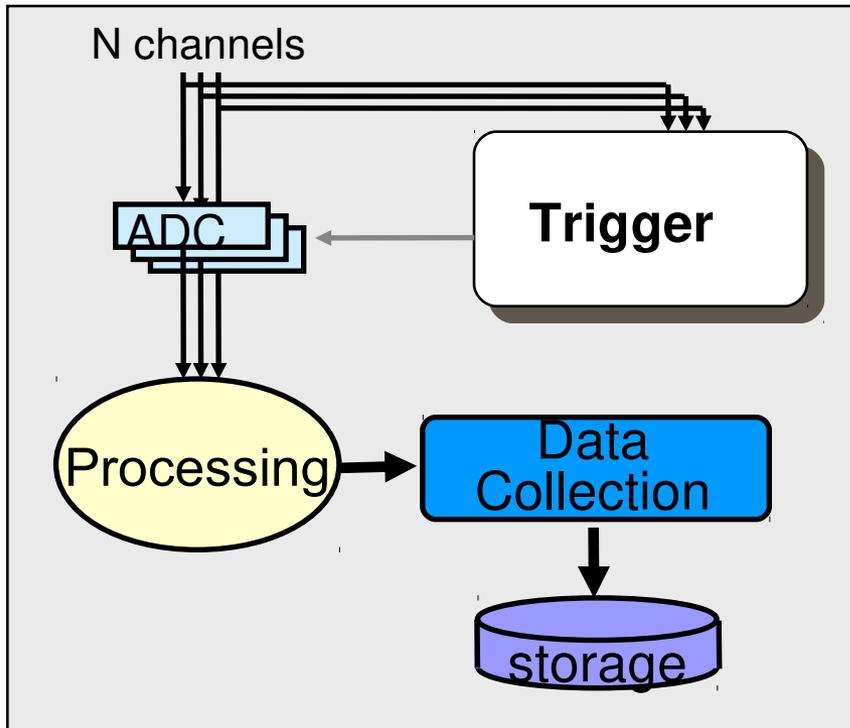
bottlenecks: PU and storage



- a single PU can be a limit
 - collect / reformat / compress data can be heavy
 - simultaneously writing storage
- final storage too:
 - VME up to 50MB/s
→ 1TB in 6h
too many disks in a week!

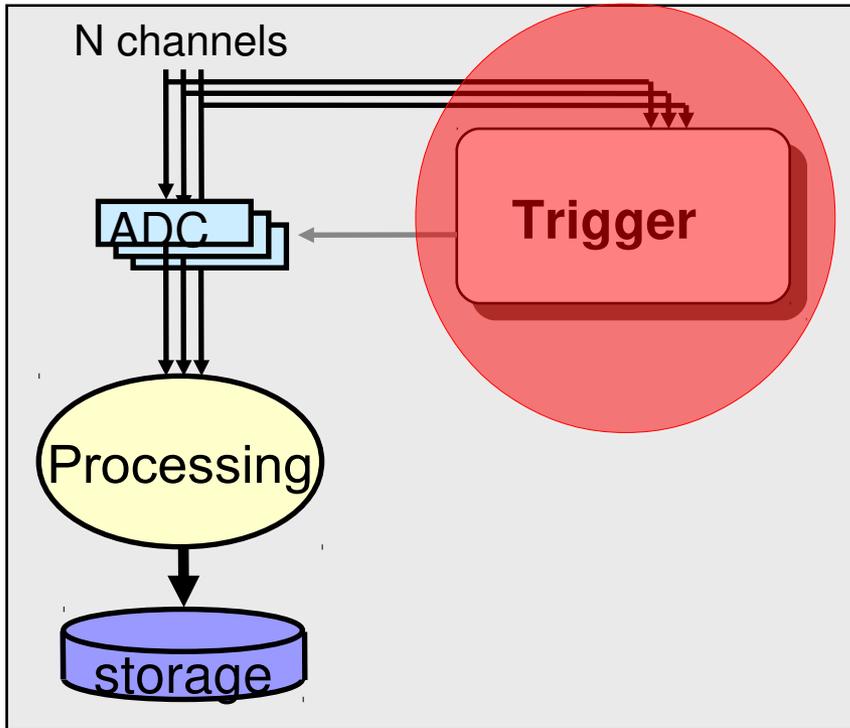
Laptop SATA disk: ~ 100 MB/s
USB2: ~ 60 MB/s

→ decouple storage from PU



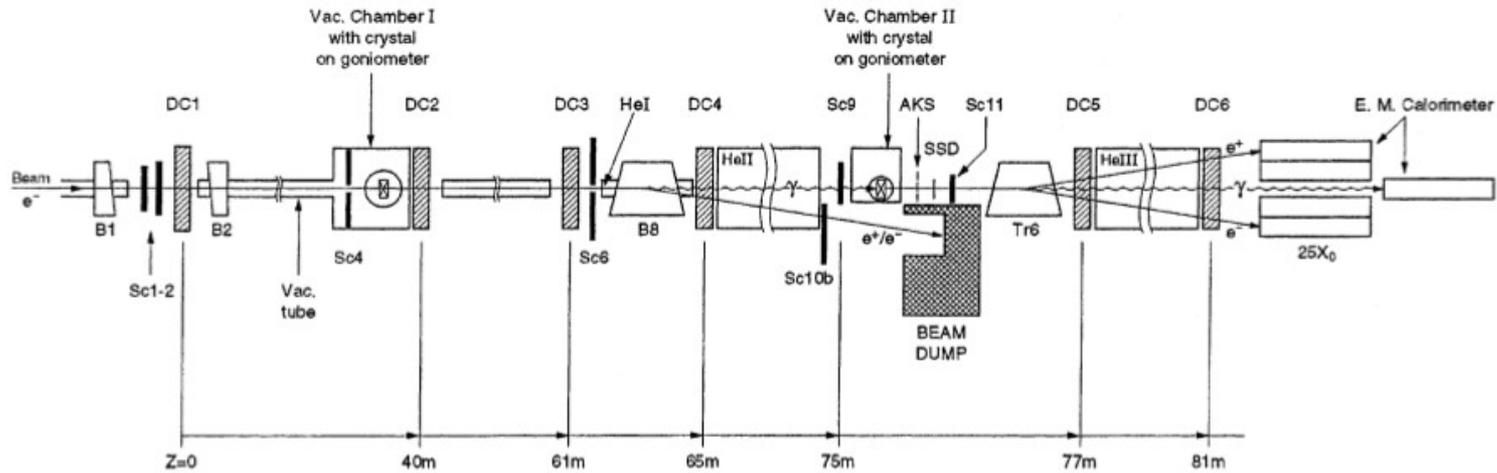
- data transfer data → dedicated “Data Collection” unit to format, compress and store
- more room for smarter processing or decreased dead time on non-buffered ADCs

bottlenecks: trigger



- to reduce data rates
(to avoid storage issues)
→ non-trivial trigger
- complexity may already hit
manageability limits for discrete
logic (latency!)
- integrated, programmable logic
came to rescue (FPGA)
→ latency may go down to
O(few ns)

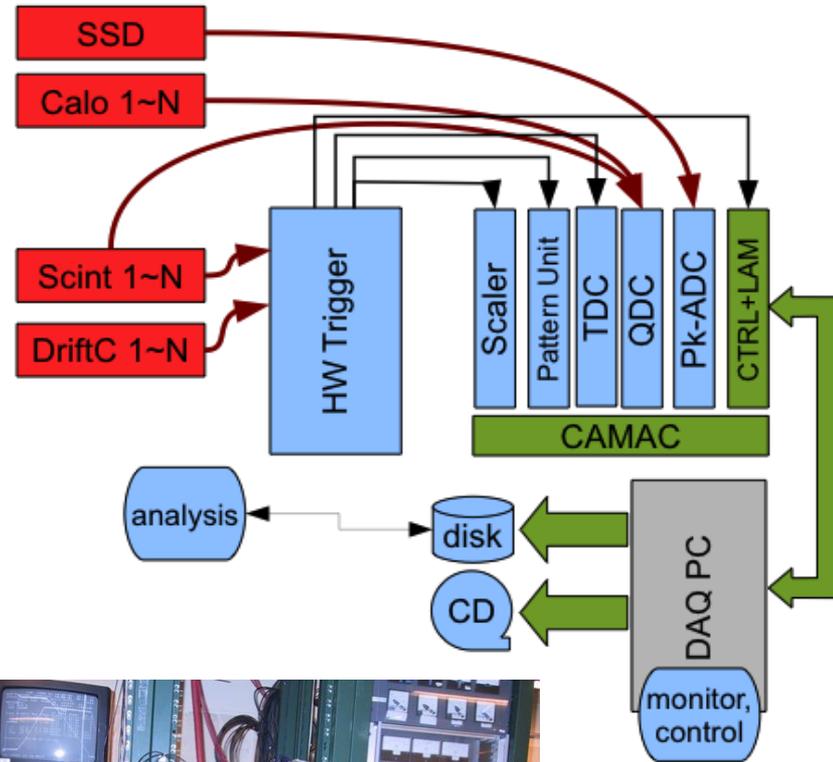
another example: NA43/63



- Radiation processes: coherent emission in crystals and structured targets, LPM suppression...
- 80/120 GeV e^- from CERN SPS slow extraction
- 2s spill every 13.5s
- Needs very high angular resolution
- Long baseline + high-res, low material detectors
→ drift Chambers
- 10 kHz limit on beam for radiation damage
→ 2-3 kHz physics trigger

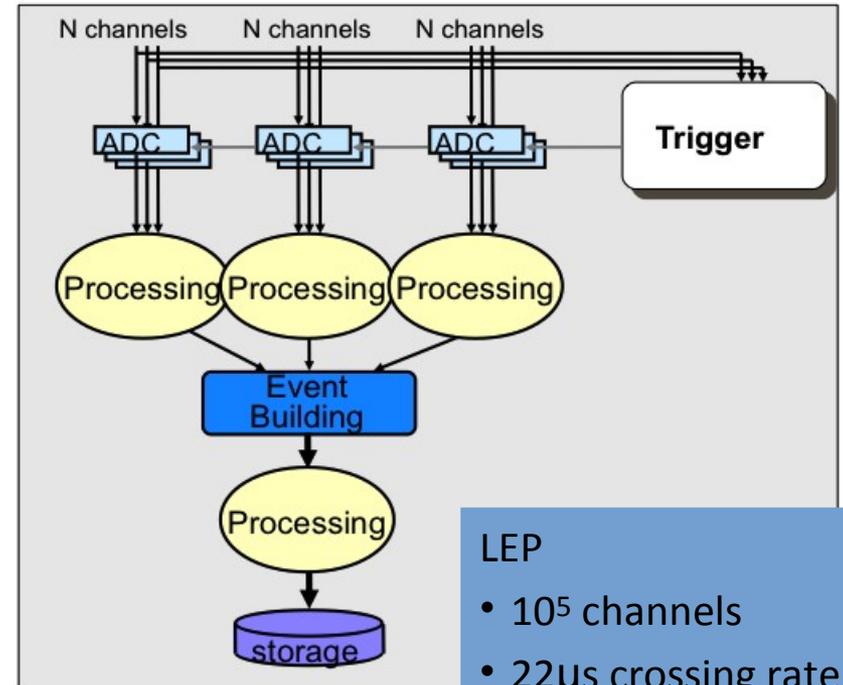
NA43/63

- 30-40 TDC, 6-16 QDC, 0-2 PADC
(depending on measurement)
- CAMAC bus
1 MB/s, no buffers, no Z.S.
- single PC readout
- NIM logic trigger
(FPGA since 2009)
 - pileup rejection
 - fixed deadtime



step three: multiple PUs (SBC)

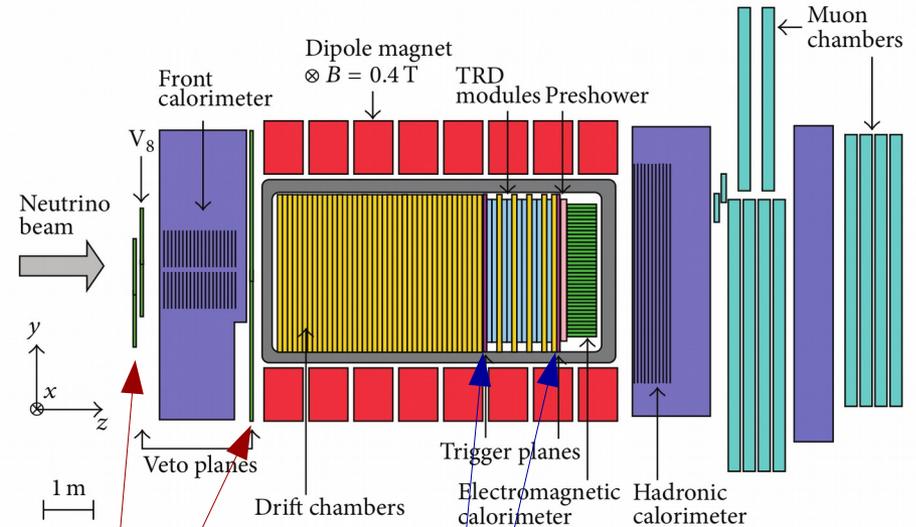
- e.g.: CERN LEP experiments
- complex detectors, moderate trigger rate, very little background
- little pileup, limited channel occupancy
- simpler, slow gas-based main trackers



- LEP
- 10^5 channels
 - $22\mu\text{s}$ crossing rate
– no event overlap
 - single interaction

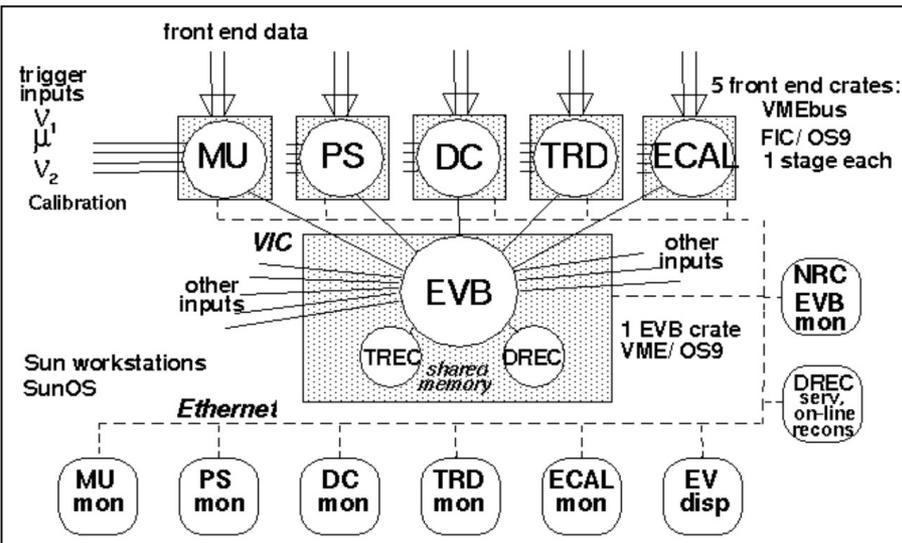
NOMAD (1995-1998)

- Search for $\nu_{\mu} \rightarrow \nu_{\tau}$ oscillations at the CERN WA neutrino facility (WANF)
- $2.4 \times 2.4 \text{ m}^2$ fiducial (beam) area
- Two 4ms spills with 1.8×10^{13} P.o.T. each (ν spills)
- One (2s) slow-extraction spill (μ spill)
- 14.4s cycle duration

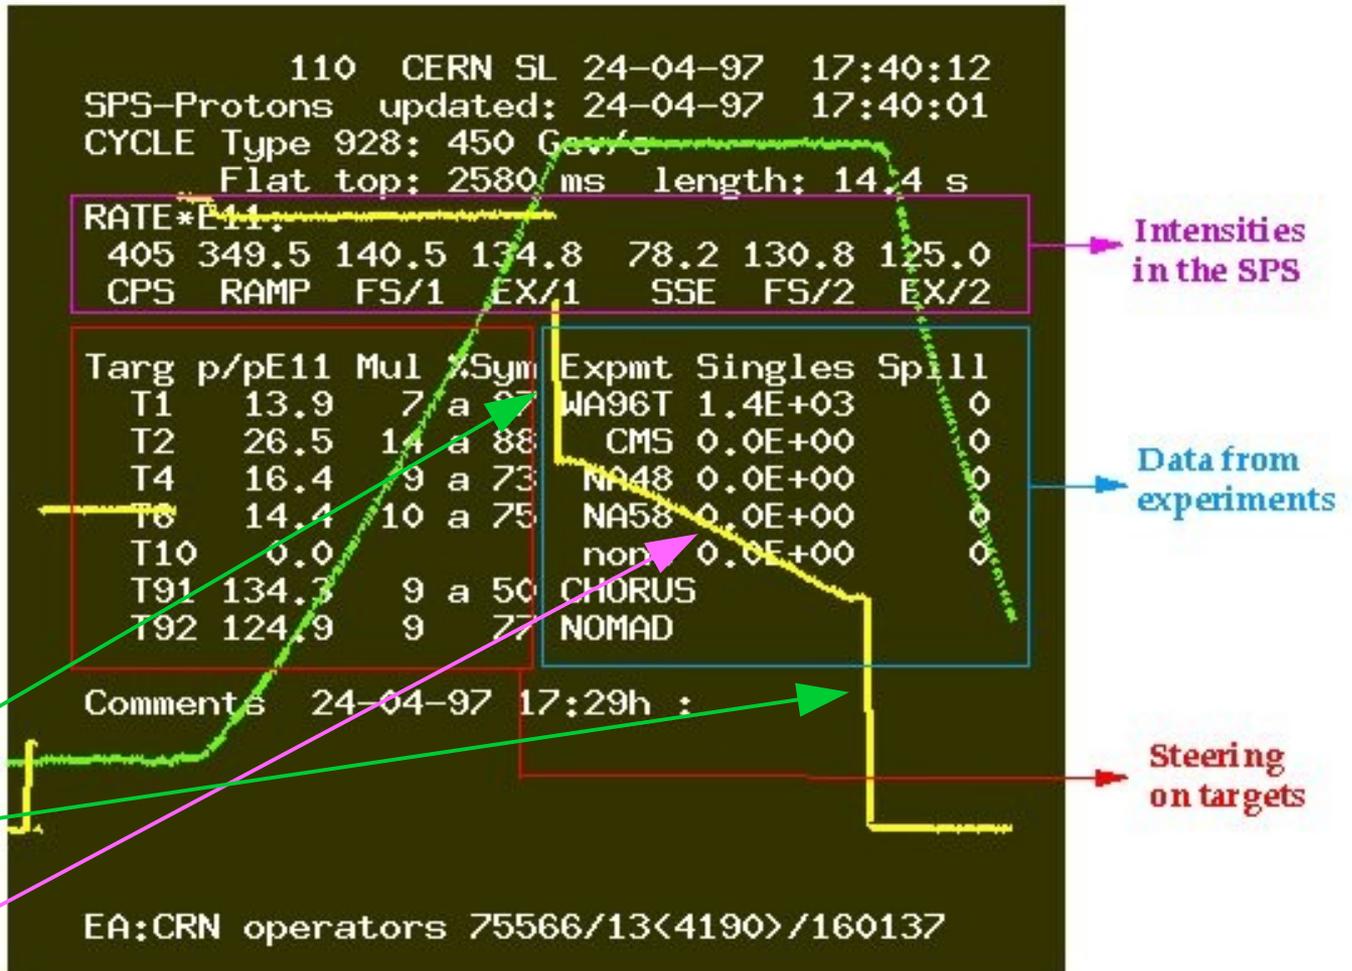


veto counters trigger counters

→ DAQ layout



WANF - SPS SuperCycle



14.4 s cycle length

2 x 4 ms neutrino spills (f/s extractions)

1 x 2 s muon spill (slow extraction)

f/s extractions

slow extraction

Intensities in the SPS

Data from experiments

Steering on targets

triggering once more ...

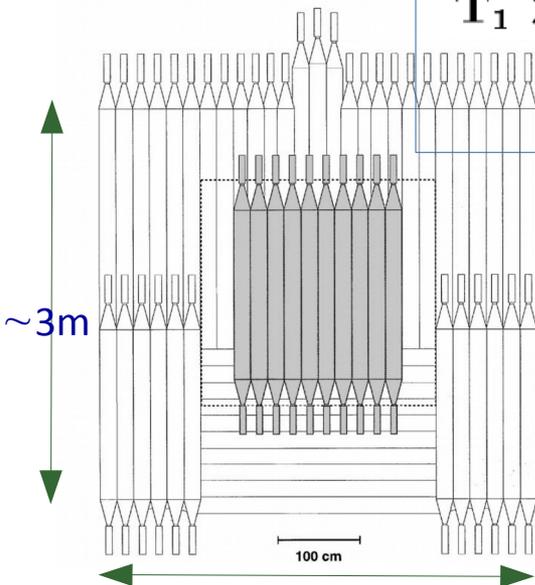
menu for NOMADs:

v-spill triggers

$\bar{V} \times T_1 \times T_2$
 $\bar{V}_8 \times FCAL$
 $\bar{V}_8 \times FCAL' \times T_1 \times T_2$
 $\overline{T_1 \times T_2} \times ECAL, \bar{V}_8 \times ECAL$
RANDOM

μ -spill triggers

$V \times T_1 \times T_2$
 $V_8 \times T_2$
 $V_8 \times T_1$
 $V_8 \times T_1 \times T_2 \times FCAL'$
 $V \times T_1 \times T_2 \times ECAL$

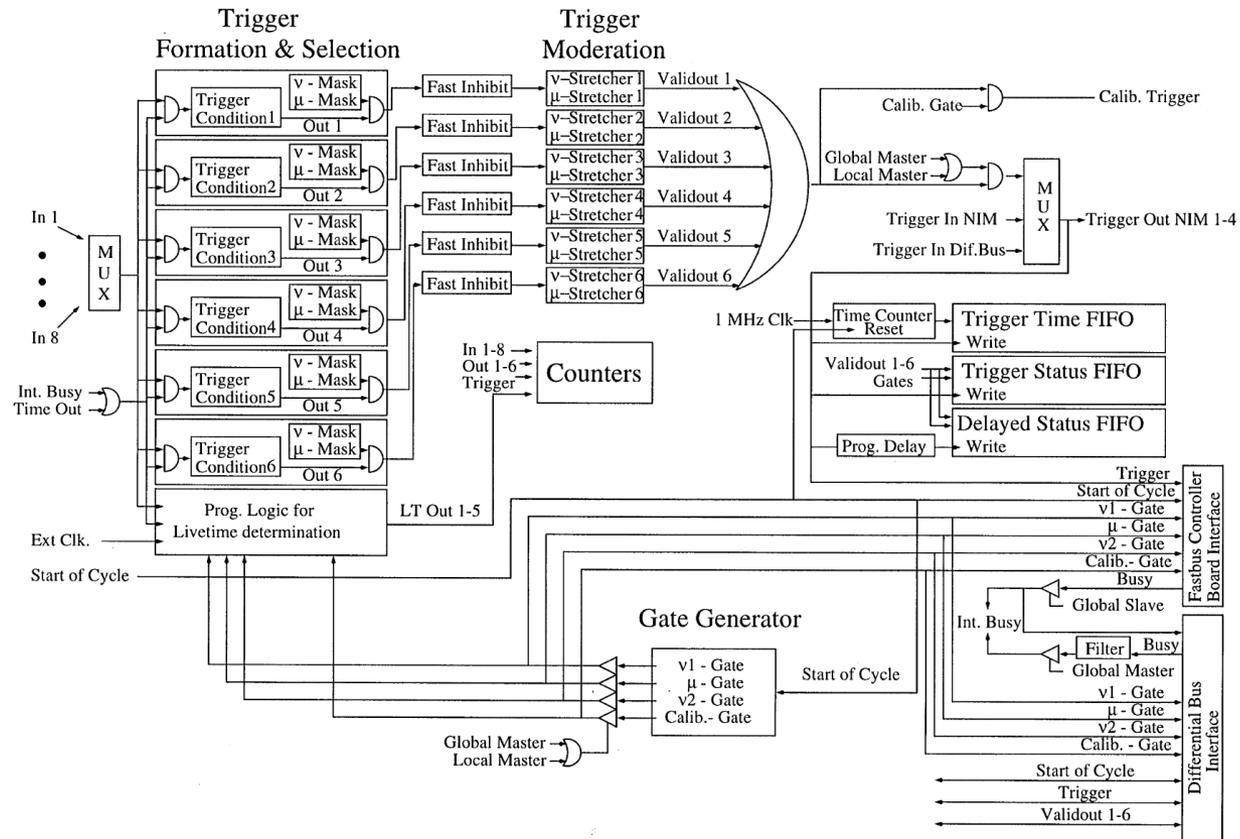


veto counters (central shaded area is V8)

triggering → FPGAs at work

MOdular TRigger for NOmad (MOTRINO):

6 VME boards providing local and global trigger generation and propagation



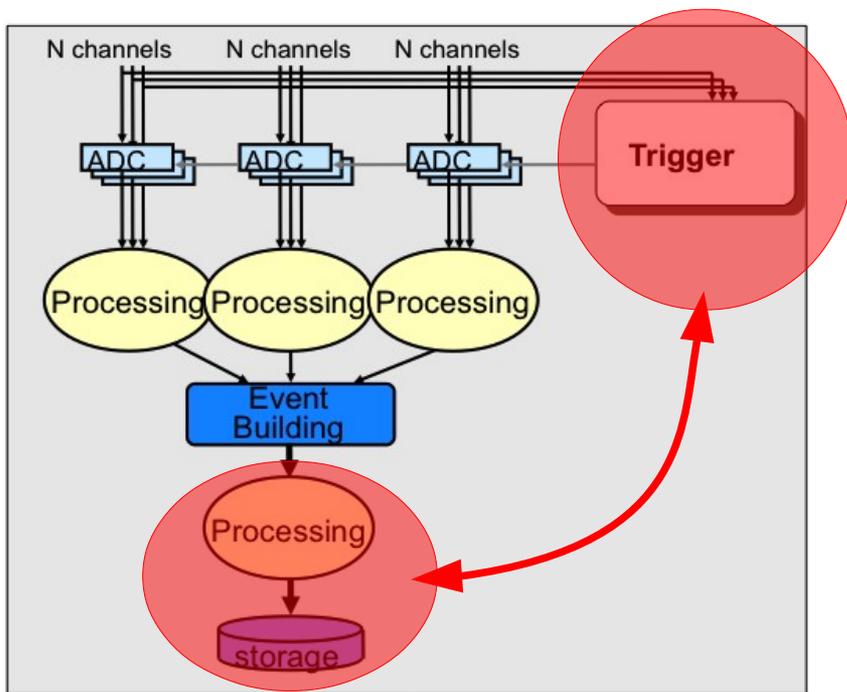
DAQ

- FASTBUS digitisers:
 - ~200 (either 64 or 96 channel) xDC boards [x=Q,P,T]
 - $O(\geq 2 \text{ us})$ conversion time, 256 event buffers
- VME readout and processing:
 - Motorola 68040 FIC8234 (OS9 real-time system) VME PUs
 - 5 for readout + 1 for event building
- Typically
 - ~4 kHz of neutrino triggers (~15 evts in each 4ms spill)
 - ~30 Hz of muon triggers (~60 evts in each 2s spill)
 - 256-events in off-spill calibration cycles (calibration triggers)

readout sequence

- On-spill on-board buffering
 - Off-spill (i.e. off-beam) data transfer and processing
 - on spill (or calibration cycle): on-board event buffering (no way to read event by event)
 - end of spill (or calibration cycle): block transfer to VME
 - then event building + storage
 - monitoring and control on SunOs and Solaris workstations
- dead time in v spills: ~10% due to digitisation

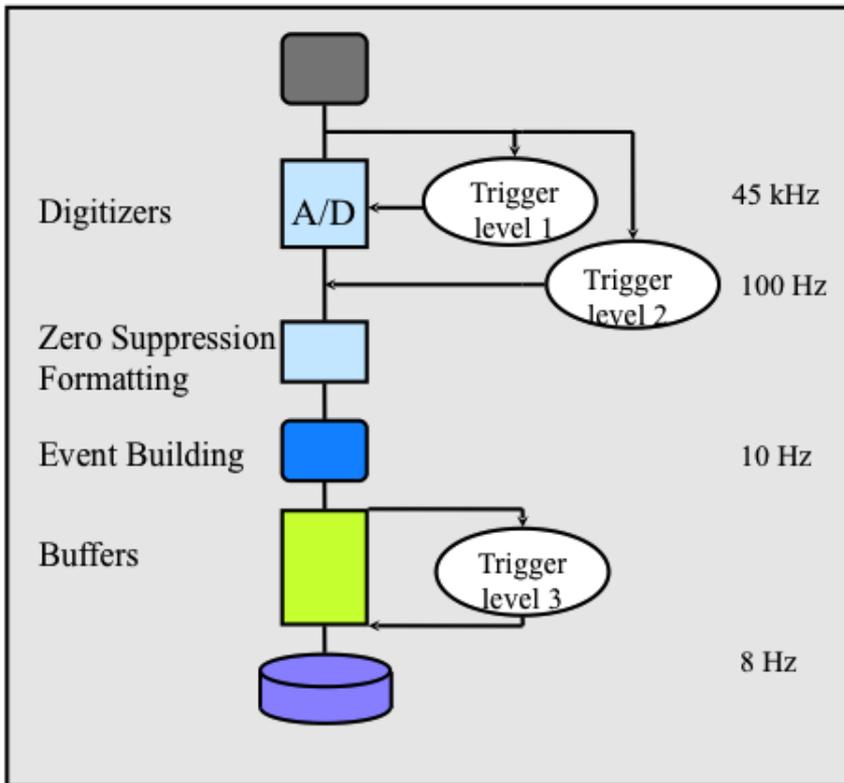
more bottlenecks ?



- trigger complexity \leftrightarrow storage
- single HW trigger not sufficient to reduce rate
- add L2 Trigger
- add HLT

step four: multi-level trigger

Typical Trigger / DAQ structure at LEP



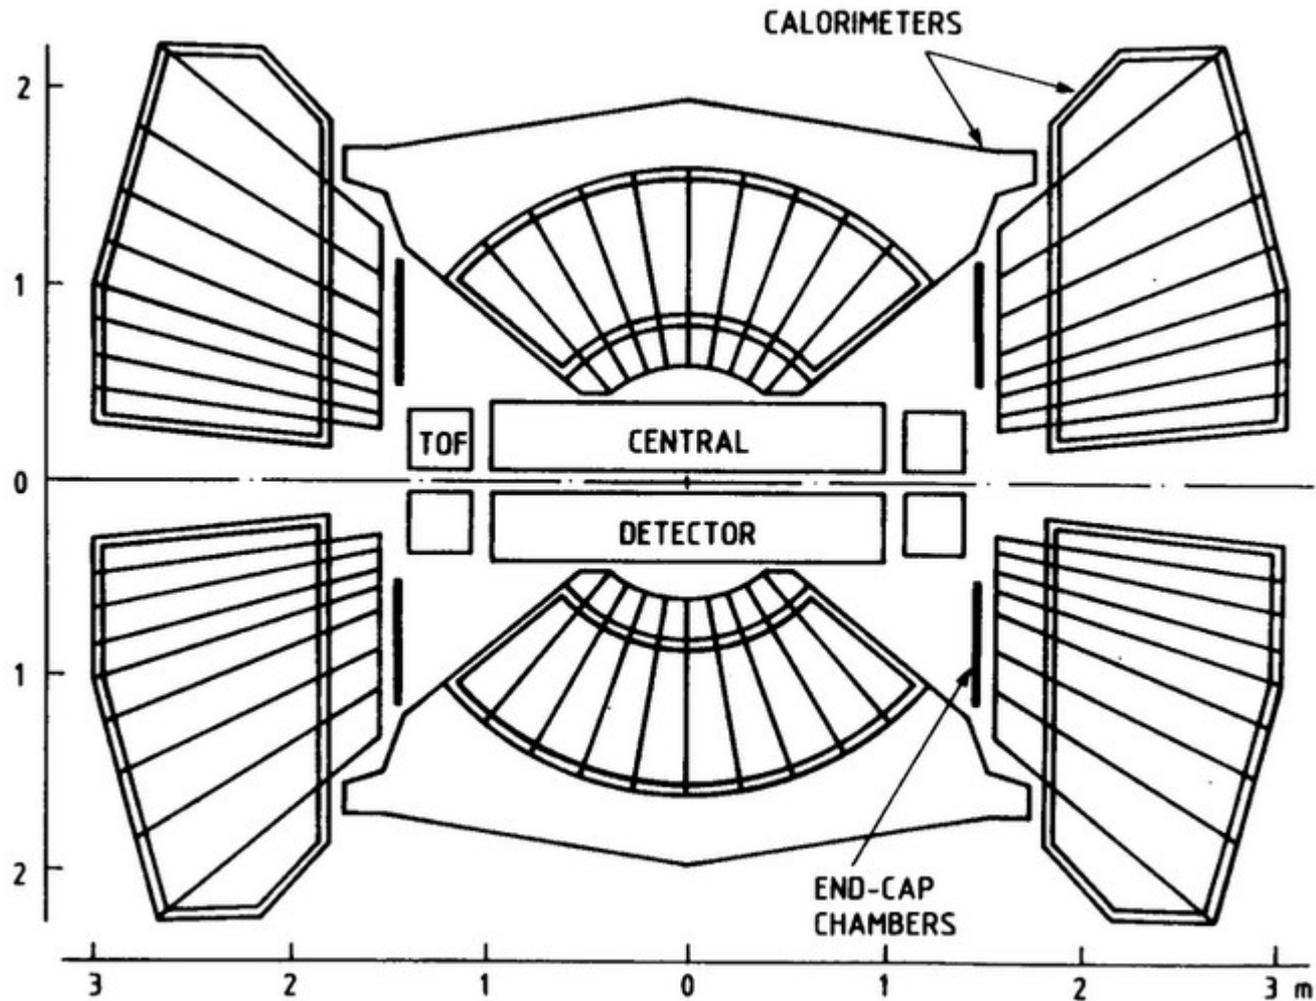
- more complex filters
 - slower
 - applied later in the chain

see Trigger lectures

LEP

- 10^5 channels
- $22\mu\text{s}$ crossing rate
 - no event overlap
- single interaction
- L1 $\sim 10^3$ Hz
- L2 $\sim 10^2$ Hz
- L3 $\sim 10^1$ Hz
- $100\text{kB/ev} \rightarrow 1\text{MB/s}$

Upgraded UA2 experiment (1988-1991)



Upgraded UA2 experiment (1987-1991)

High-lumi pp collisions @ CERN SppS:

$$\sqrt{s} = 630 \text{ GeV}$$

$$L = 5 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1} \text{ (one order of magnitude increase)}$$

Goal:

W/Z physics

QCD

top quark and SUSY particle discovery

→ robust theoretical prediction for this

Complex trigger signatures:

em, jet and missing P_T

Upgraded UA2 experiment (1987-1991)

Three-level trigger selection:

L1 from on detector hardware

L2 over a dedicated processor

L3 over FASTBUS processors (ALEPH event builder)

DAQ readout & monitoring:

CAMAC & FASTBUS → VAX/VMS platforms

Upgraded UA2 experiment (1987-1991)

Three-level trigger selection:

L1 from on detector hardware

L2 over a dedicated processor

L3 over FASTBUS processors (ALEPH event builder)

DAQ readout & monitoring:

CAMAC & FASTBUS → VAX/VMS platforms

→ theorist were right but ...

Upgraded UA2 experiment (1987-1991)

Three-level trigger selection:

L1 from on detector hardware

L2 over a dedicated processor

L3 over FASTBUS processors (ALEPH event builder)

DAQ readout & monitoring:

CAMAC & FASTBUS → VAX/VMS platforms

→ theorist were right but ...

unfortunately nature was wrong!

Upgraded UA2 experiment (1987-1991)

Three-level trigger selection:

L1 from on detector hardware

L2 over a dedicated processor

L3 over FASTBUS processors (ALEPH event builder)

DAQ readout & monitoring:

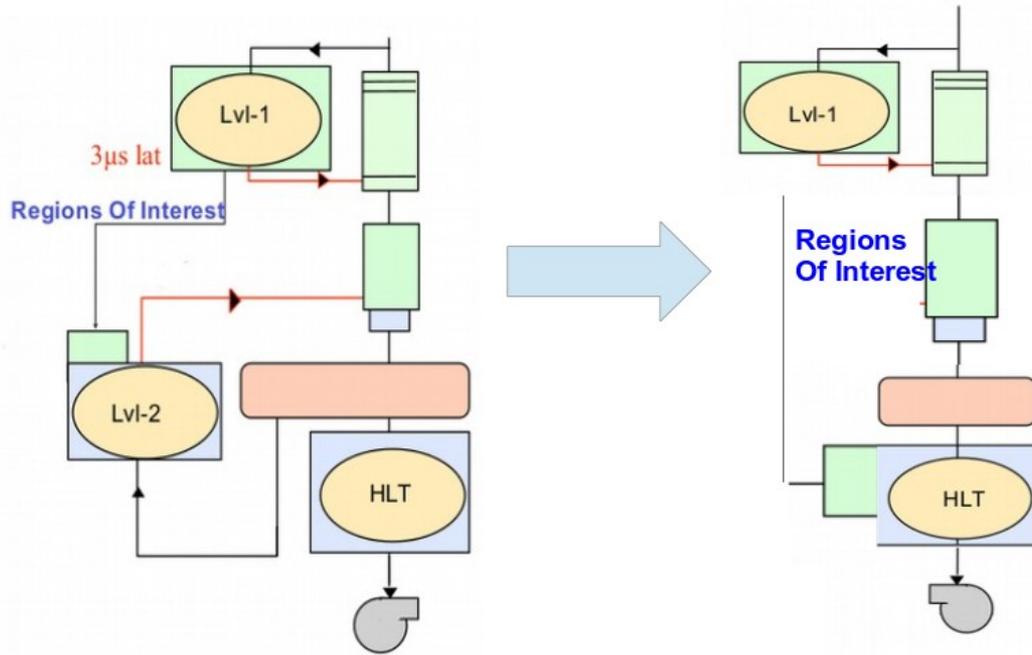
CAMAC & FASTBUS → VAX/VMS platforms

→ theorist were right but ...

unfortunately nature was wrong!

Nevertheless many new/better measurements and observations of SM processes

ATLAS (from run-1 to run-2)



→ Merge L2 and L3 into a single HLT farm

- preserve Region of Interest but dilute the farm separation and fragmentation
- increase flexibly, computing power efficiency

off-topic: event-selection latency

Typical (ATLAS) numbers:

- L1 : $O(1 \mu\text{s in real-time}) \rightarrow \text{let say} = 2.5 \mu\text{s}$
- L2 : $O(10 \text{ ms}) \rightarrow \text{let say} = 40 \text{ ms}$
- L3(HLT) : $O(\text{s}) \rightarrow \text{let say} = 1 \text{ s}$

Q: do the 3 numbers mean the same thing ?

latency and real-time

real time: system must respond within some fixed delay

→ *Latency = Max Latency*

→ over fluctuations bad, will create dead time

non-real-time: system responds as soon as it's available

→ *Latency = Mean Latency*

→ over fluctuations fine, shouldn't create dead time

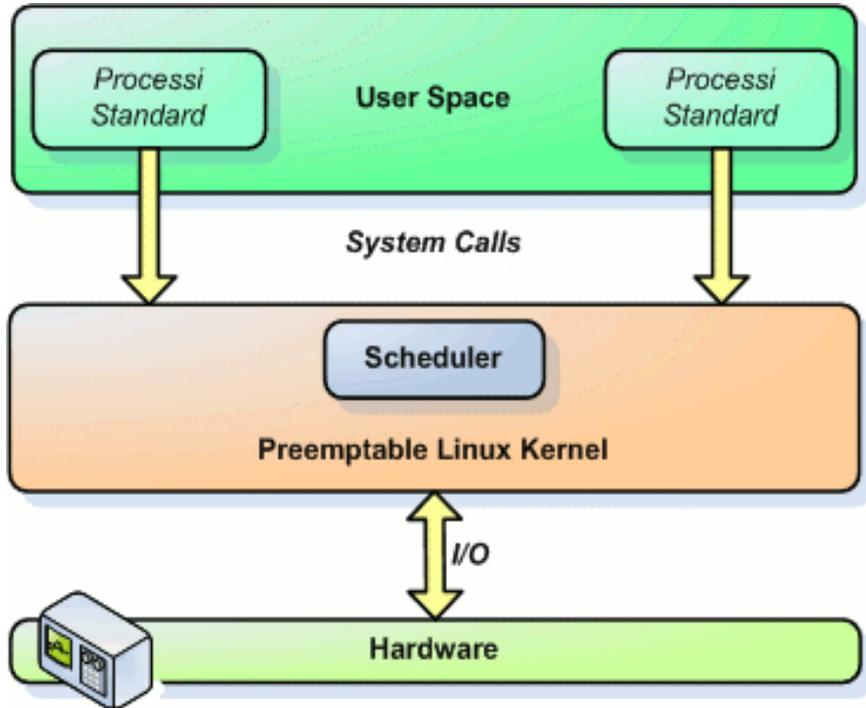
real time o.s. :

very stable time delay in responding to events

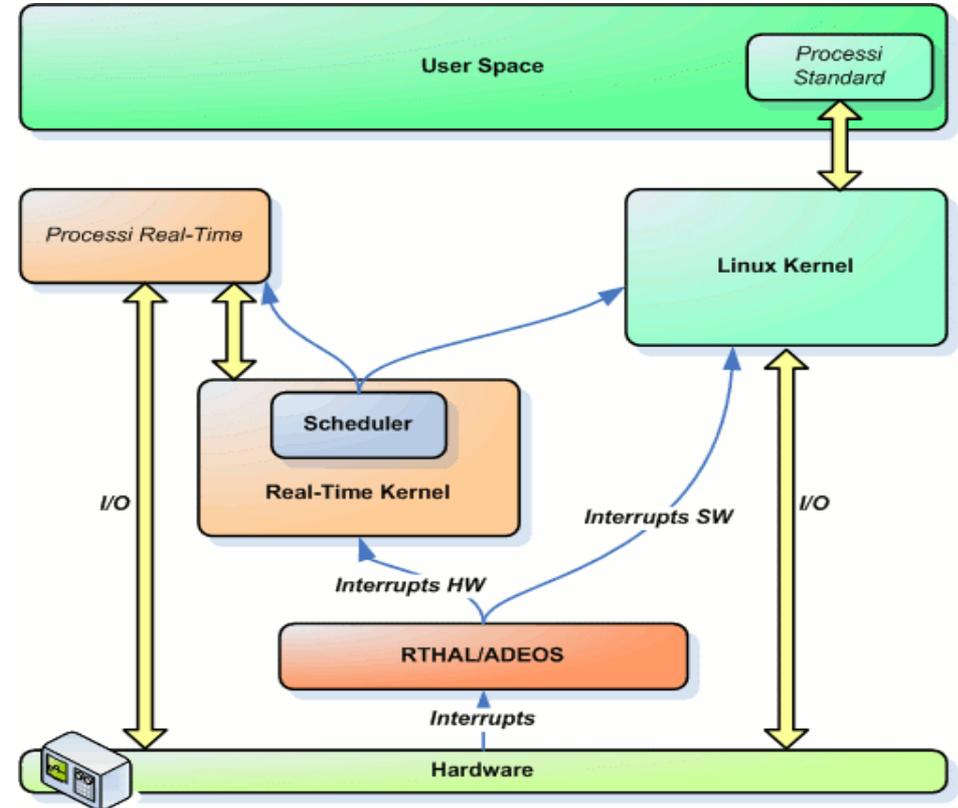
standard unix kernels are not real time:

a system call can in principle take any time

off-topic: real-time linux

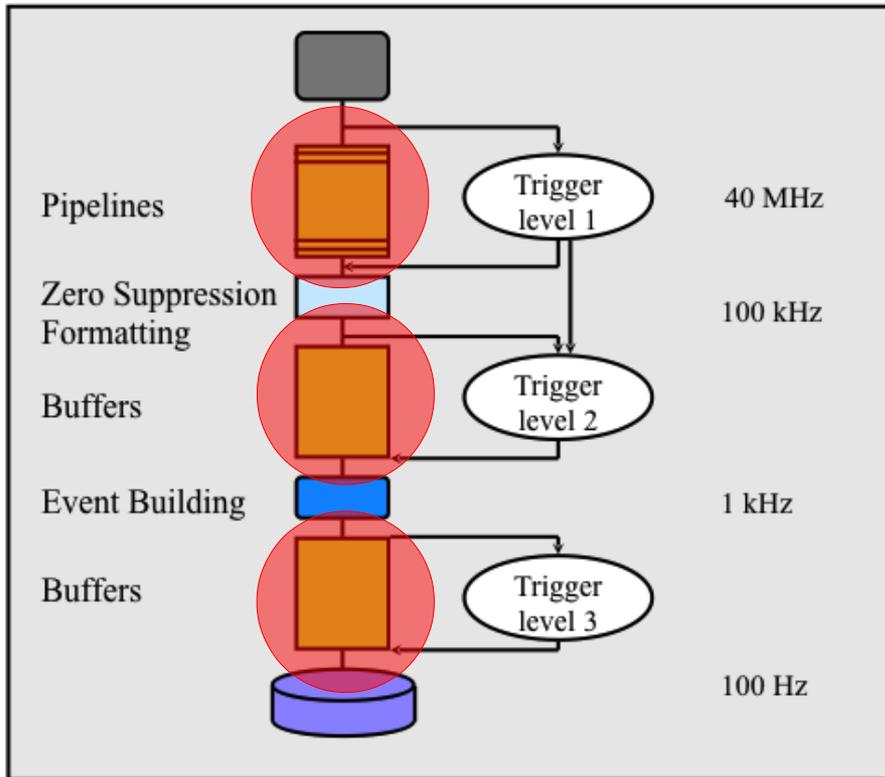


Low-latency patch
(Ubuntu Studio):
Interruptible linux kernel



RTAI: linux kernel runs as a
higher priority application

step five: dataflow control

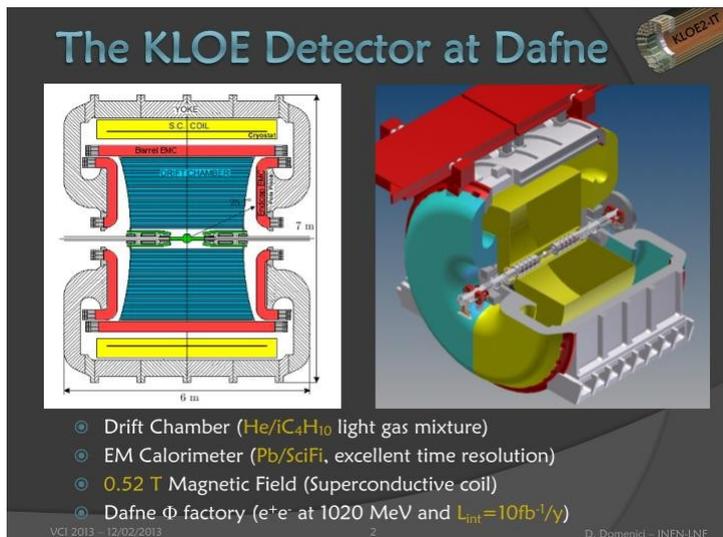


- Buffers are not the <final solution> they can overflow due to:
 - bursts
 - unusual event sizes
- Discard
 - local, or
 - “backpressure”, tells lower levels to discard

Who controls the flow?
FE (*push*) or EB (*pull*)

a *push* example: KLOE

- DAΦNE e^+e^- collider in Frascati
- CP violation parameters in the Kaon system
- “factory”: rare events in a high-rate beam
- 10^5 channels
- 2.7 ns crossing rate
 - rarely event overlap
 - “double hit” rejection
- high rate of small events
- L1 $\sim 10^4$ Hz
 - $2\mu\text{s}$ fixed dead time
- HLT $\sim 10^4$ Hz
 - \sim COTS, cosmic rejection only
- 5 kB/ev \rightarrow 50 MB/s [design]

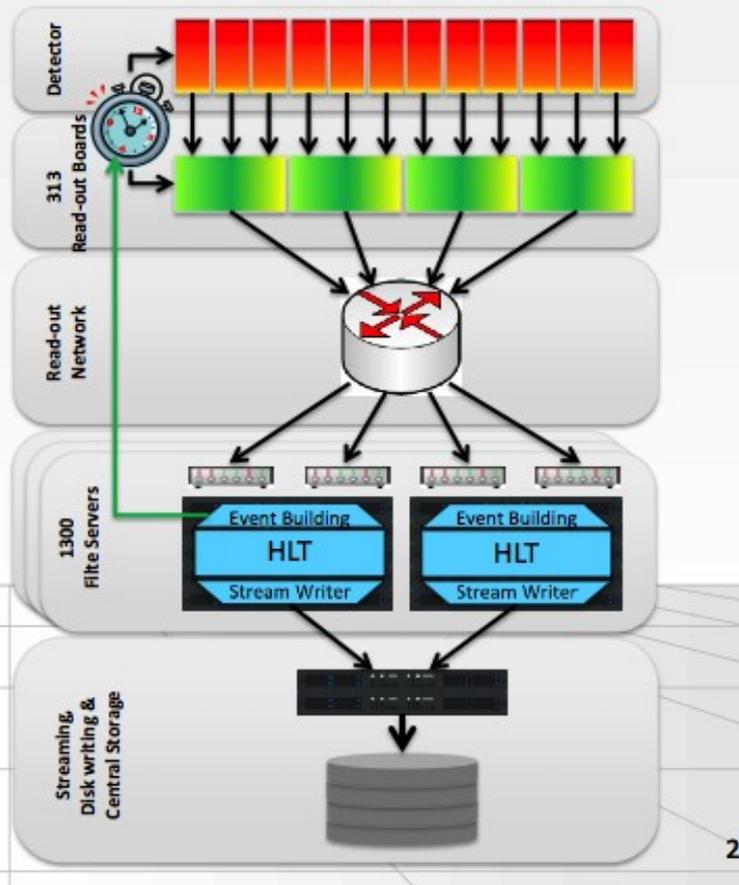


LHCb: network is dataflow



From Front-End to Hard Disk

- $O(10^6)$ Front-end channels
- 300 Read-out Boards with 4 x 1 Gbit/s network links
- 1 Gbit/s based Read-out network
- 1500 Farm PCs
- >5000 UTP Cat 6 links
- 1 MHz read-out rate
- Data is pushed to the Event Building layer. There is no re-send in case of loss
- Credit based load balancing and throttling



The LHCb Data Acquisition during LHC Run 1
CHEP 2013

more info in "TDAQ for the LHC experiments"

looking forward to LS2 and beyond

On some long term, all experiments looking forward to significant increase in L1 trigger rate and bandwidth.
ALICE and LHCb will pioneer this path during LS2

DAQ@LHC Workshop

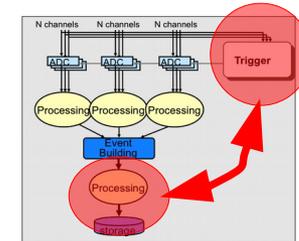
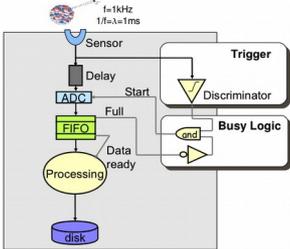


- First level trigger for Pb-Pb interactions **500 Hz** → **50 kHz**
- 22 MB/event
 - **1 TB/s readout** → **500 PB/month**
- Data volume reduction
 - on-line full reconstruction
 - discard raw-data
- Combined DAQ/HLT/offline farm
 - COTS, FPGA and GPGPU



- **1 MHz** → **40 MHz** readout and event building → trigger-less
 - trigger support for staged computing power deployment
- 100 kB/event
 - on-detector zero suppression → rad-hard FPGA
 - 4 TB/s event-building

trends



- Integrate synchronous, low latency in front end
 - limitations do not disappear, but decouple (factorise)
 - all-HW implementation
 - isolated in replaceable(?) components
- Use networks as soon as possible
- Deal with dataflow instead of latency
- Use COTS network and processing
- Use “network” design already at small scale
 - easily get high performance with commercial components

take care, lot of issues not covered:

Hw configuration

Sw configuration

Hw control & recovery

Sw control & recovery

Monitoring

...

Thank you for your patience ...

Lost & Found (off-topics)

Appendix A: Cables and Transmission Lines

Spoken about signals, amp.s, digitisers, ... but ...

... almost nothing about how signals are transmitted over long distances. ***Is there any issue ?***

Q(1): what is a cable (for a single signal) ?

a couple of ideal conductors ($R=C=L=0$) ?

Q(2): which speed can it reach ?

Q(3): what's its impedance ?

Q(4): what does it do to your signal ?

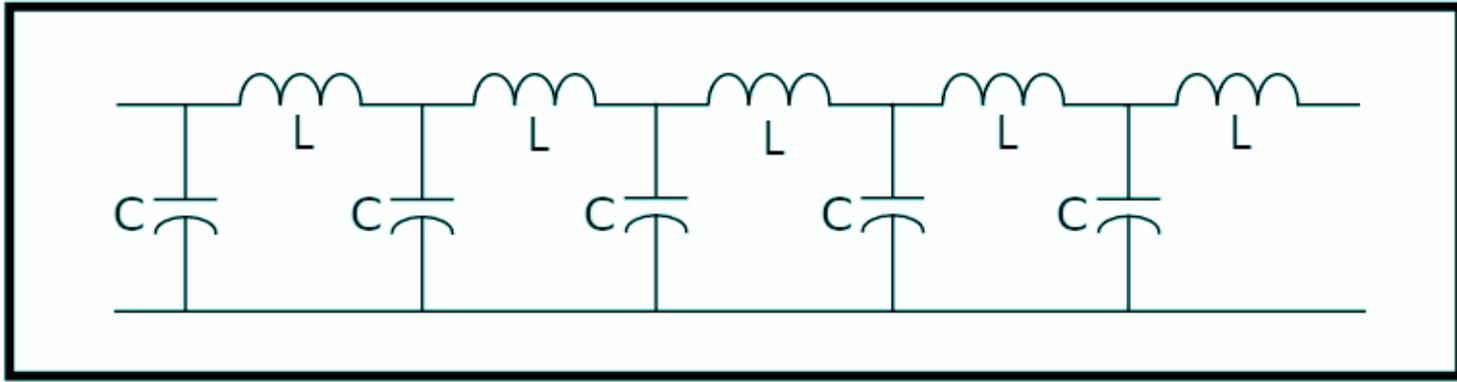
Ok the full line must be properly matched:

$$Z(\text{out}) = Z(\text{cable}) = Z(\text{in})$$

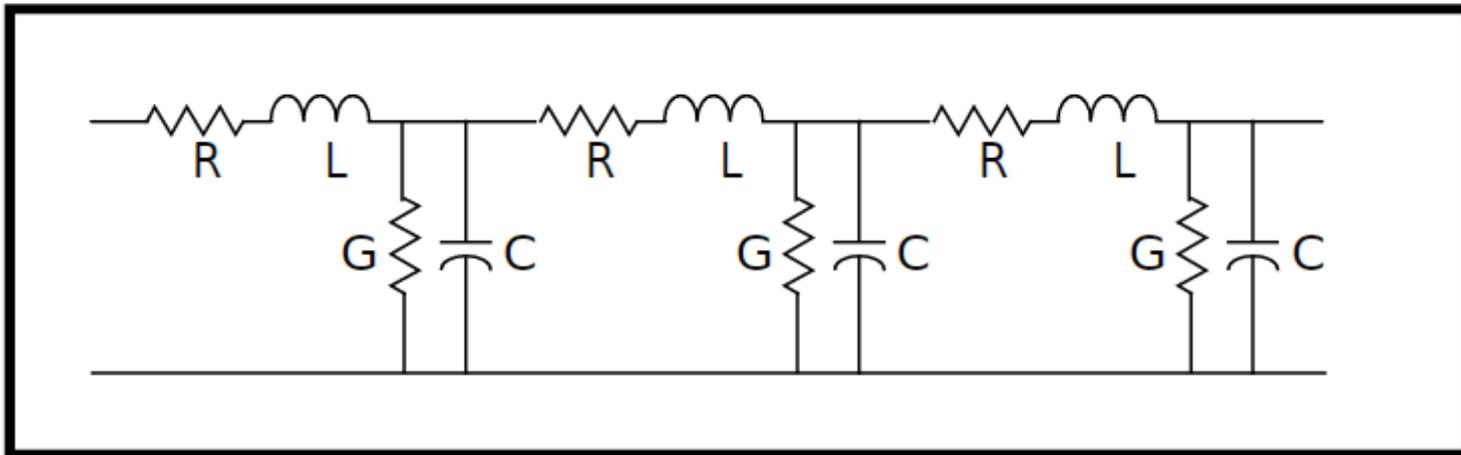
That's all ?

Cables and Transmission Lines

Lossless transmission line:



Lossy transmission line:

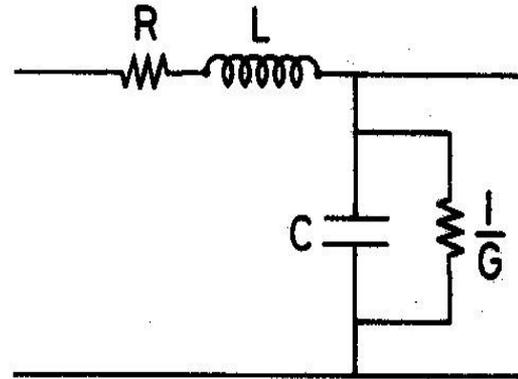


Cables

Cable element (dz):

$$L \approx \frac{\mu}{2\pi} \ln\left(\frac{b}{a}\right) \quad [\text{H/m}]$$

$$C \approx \frac{2\pi\epsilon}{\ln(b/a)} \quad [\text{F/m}]$$



R depends on the frequency (skin effect)

G should be negligible

$$Z = (L/C)^{1/2}$$

$$v_p = (LC)^{-1/2} = (\mu\epsilon)^{-1/2}$$

Cables

Equation for standing waves:

$$\frac{\partial^2 V}{\partial z^2} = LC \frac{\partial^2 V}{\partial t^2} + (LG + RC) \frac{\partial V}{\partial t} + RGV$$

solution:

$$\frac{d^2 V}{dz^2} = (R + i\omega L)(G + i\omega C)V = \gamma^2 V$$

$$\gamma = \alpha + ik = \sqrt{(R + i\omega L)(G + i\omega C)}$$

R usually dominated by the skin effect:

$$R(\omega) = r * D / (4 * \delta)$$

r = resistance per unit length

D = diameter internal conductor

δ = skin depth $\sim 1/\sqrt{\omega}$

Cable Losses

Neglecting the transconductance G :

$$\alpha = R(\omega)/(2Z_0) \sim c\sqrt{\omega}$$

$$k = \omega\sqrt{RC} = \omega/(\beta c)$$

$$V(z, t) = V_1 \exp(-\alpha z) \exp[i(\omega t - kz)]$$

50-Ohm fast ($v = 4$ ns/m) CERN-store cables:

04.61.11.F - COAXIAL CABLE 50 OHM - TYPE C-50-6-1

04.61.11.H - COAXIAL CABLE 50 OHM - LOW LOSS - TYPE C-50-11-1

$f(-3\text{db}, 40 \text{ m, cable C-50-6-1}) \sim 120 \text{ MHz}$

$f(-3\text{dB}, 40 \text{ m, low loss cable}) \sim 640 \text{ MHz}$

Signal Distortions

Time parameter:

$$\alpha \sim \mu \sqrt{f}$$
$$\tau_0 = (\mu z)^2 / \pi$$

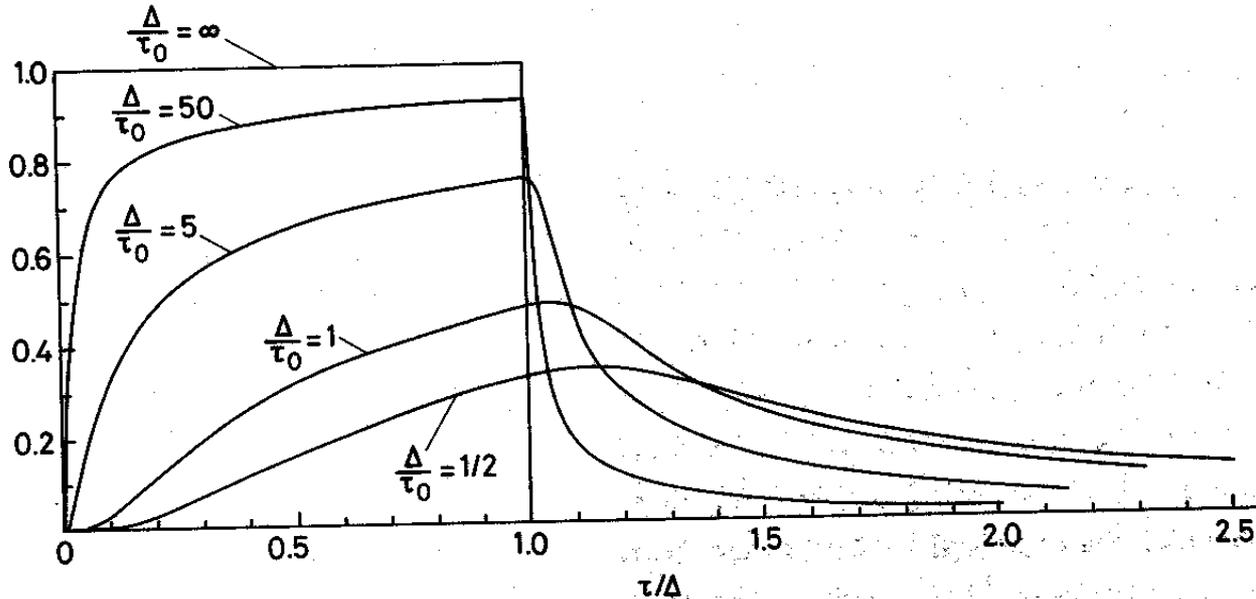
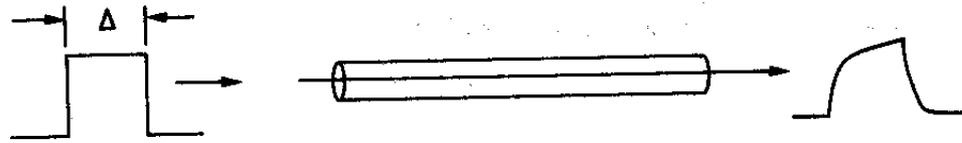
$\mu z \sim 32 \cdot E-6$ (C-50-6-1), $14E-6$ (low loss cables)

$$\tau_0 \sim 320 \text{ ns} \text{ (C-50-6-1)}$$

$$\tau_0 \sim 60 \text{ ns} \text{ (low loss cables)}$$

*** Take care: would like $\tau_0 \ll \tau(\text{signal})$

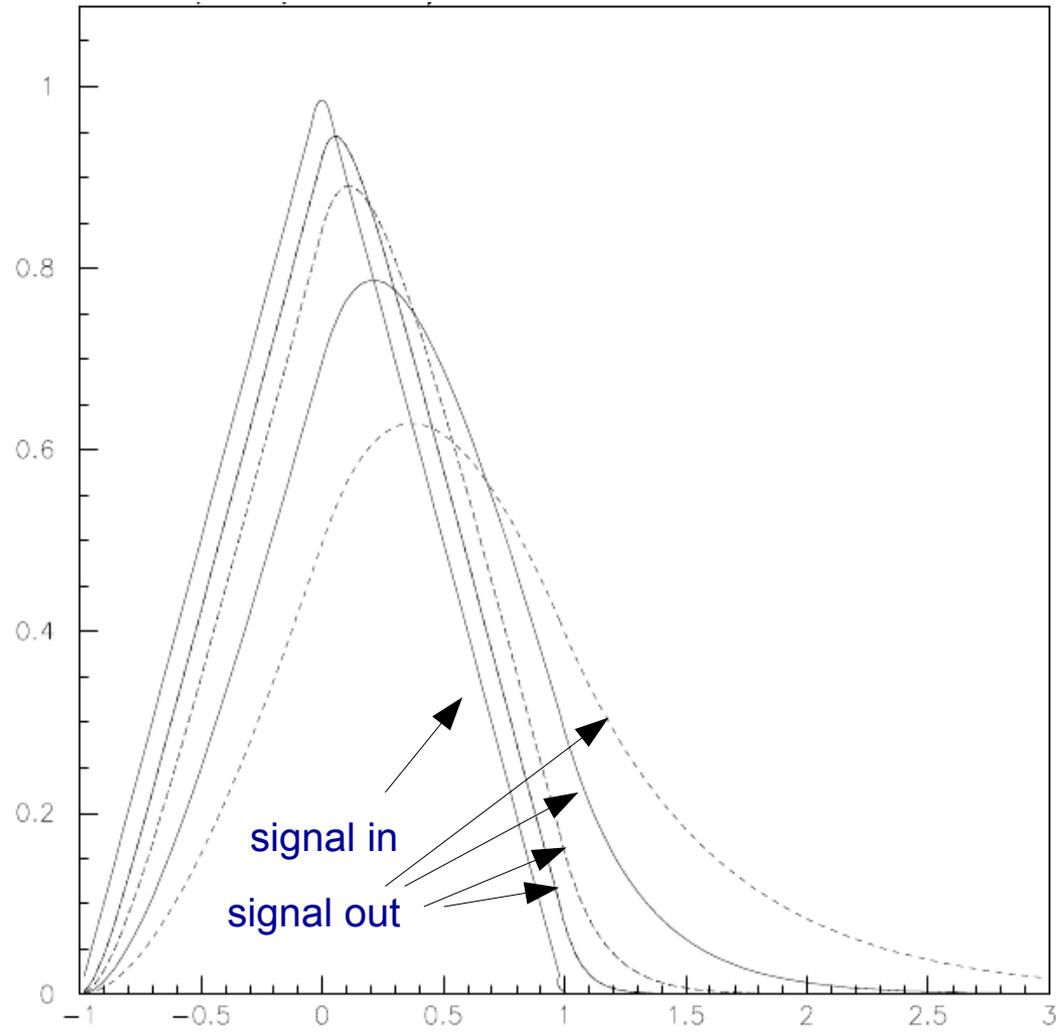
Digital Pulse Distortions



$$\text{for } \alpha \sim a\sqrt{f}$$
$$\tau_0 = (az)^2 / \pi$$

Bandwidth Effects – Analog Signals

~1ns analog-signal response for
BW ~ 300, 150, 75, ... MHz



Appendix B: backtrace

Segfaulting ? Have a look at backtrace:

https://www.gnu.org/software/libc/manual/html_node/Backtraces.html

BACKTRACE(3)

Linux Programmer's Manual

BACKTRACE(3)

NAME

backtrace, backtrace_symbols, backtrace_symbols_fd - support for application self-debugging

SYNOPSIS

```
#include <execinfo.h>
```

```
int backtrace(void **buffer, int size);
```

```
char **backtrace_symbols(void *const *buffer, int size);
```

```
void backtrace_symbols_fd(void *const *buffer, int size, int fd);
```

HowTo

1) file “my_segf.cxx” : install a signal handler to print the backtrace

```
#include <stdio.h>
#include <execinfo.h>
#include <signal.h>
#include <stdlib.h>
#include <unistd.h>
```

```
void handler(int sig) {
    void *array[10];
    size_t size;

    // get void*'s for all entries on the stack
    size = backtrace(array, 10);

    // print out all the frames to stderr
    fprintf(stderr, "Error: signal %d:\n", sig);
    backtrace_symbols_fd(array, size, STDERR_FILENO);
    exit(1);
}
```

```
void baz() {
    int *foo = (int*)-1; // make a bad pointer
    printf("%d\n", *foo); // causes segfault
}

void bar() { baz(); }
void foo() { bar(); }

int main(int argc, char **argv) {
    signal(SIGSEGV, handler); // install our handler
    foo(); // this will call foo, bar, and baz. Baz segfaults.
}
```

2) compile with -g debug flag on:

```
g++ -g -rdynamic my_segf.cxx -o  
my_segf
```

3) get the crash:

```
[roberto@bob-laptop ~]$ ./my_segf  
Error: signal 11:  
./my_segf(_Z7handleri+0x1c)[0x400a52] ← handler  
/lib64/libc.so.6(+0x347e0)[0x7fa55f1c07e0] ← libc  
./my_segf(_Z3bazv+0x14)[0x400aab] ← my crash  
./my_segf(_Z3barv+0x9)[0x400aca]  
./my_segf(_Z3foov+0x9)[0x400ad6]  
./my_segf(main+0x23)[0x400afc]  
/lib64/libc.so.6(__libc_start_main+0xf1)  
[0x7fa55f1ac731]  
./my_segf(_start+0x29)[0x400969]
```

4) crash is at (`_Z3bazv+0x14`) ... the function name is “`_Z3bazv`” (c++ function name mangling). How to get it ?

5) Demangle it thanks to: <http://>

6) Take the Answer: `baz()` → demangler.com/ crash is at (`baz+0x14`)

7) crash is at (baz+0x14) ... open the debugger: `gdb my_segf`

```
(gdb) info address baz
```

```
Symbol "baz()" is a function at address 0x400a55.
```

8) so crash is at address (0x499a55+0x14) ... then:

```
(gdb) info line *(0x400a55+0x14)
```

```
Line 24 of "my_segf.cxx" starts at address 0x400a65 <baz()+16>  
and ends at 0x400a7c <baz()+39>.
```

9) got it ! That's not yet the reason but ...

Appendix C: Profiling

Take care: optimize your code – first of all - where it really needs. To get it, you may use of profiling.

for C/C++ code, look (for example) at this gprof tutorial:

<http://www.thegeekstuff.com/2012/08/gprof-tutorial/>

Very simple, at least for standalone code ...