DAQ, Trigger and Offline in Future Upgrade
(more questions than answers)

Giovanni Punzi

University of Pisa and INFN

LHCB-IT meeting

Frascati, 13 Oct 2015

Introduction

- In the TTFU, we have considered and discussed several aspects of what will be needed for a Very High Lum upgrade (FU)
- The issue of "Data Handling" (DAQ/Trigger/Offline) was maybe the least developed
- There are indeed more questions than answers
- A lot of DH issues revolve around *tracking*, and tracking for a FU is still very much an open subject
- Forward tracking in particular
- Will do my best to cover the issues, drawing a lot from work I started in 2014 about FCC and XFX ideas.

Starting from the end: Physics objects

- By this I mean the high-level reconstruction of the decay of interest to its fullest extent
- Losing efficiency with luminosity defeats the purpose of a FU

→ At $10^{34 \div 35}$ we will be collecting O(25-250x) what we collect today. Say 10 kHZ → 250-2500 kHz. (LHCbU=20÷100kHz)

→ This cannot be conceivably processed (and stored) offline.

- Most event are signals, so reducing event rate is hardly an option
  - Arguments have been put forth for rare decays only
  - But it seems hard to push a FU unless based on WIDE program

→ *We must reduce event size by factors* O(25-250x)

→ From 100kB to <u>< 1kb</u> per "event"

→ *We will need to store "decay candidates" rather than "events"*

Starting from the end: Physics Objects

We will need to store "decay candidates" rather than "events"

→ Will need to perform most of the analysis in **real time**

We are already moving in this direction in Run-2 with the TURBO stream.
→ An indication that the direction is a necessity (as apposed to an approach advocated until recently, of "taking decisions as late as possible")

Doing it right is though – remember that a FU only makes sense if we can have extra-tiny systematics (in some cases this means 10^{-5}), and up to know we only made *very simple* measurements in this way (resolution O(100%))

→*testing our ability to do precision measurements with a TURBO stream will be critical to establish the feasibility of a Future Upgrade*

LHCb vs GPD processing requirements

- Data produced: LHCb-U: 5 TB/s GPDs: >150 TB/s
  - GPD solution: trigger at low-level to reduce rate to ~5 TB/s level before moving data off-detector
    - *L1 track triggers* in development at ATLAS and CMS
    - Filter Hi-PT tracks + Trigger events at low-level
    - Need reduction not only for readout, but for *processing* cost
  - FU means going to GPD-like flows at LHCb

    ...but LHCb cannot do low-level event selection, as most event contain some useful decay

    - Makes FU **much harder** than at GPD – almost FCC !

→ *LHCb FU can only readout* ***partial*** *events*

- Not only we won't be able to store full events - Now I am saying we need to make choices even *before readout*

→ *We must prepare for "detector-embedded reconstruction"*

Hardest issue for a FU is TRACKING

(hadron PID may be even harder – but will not touch on it today)

- The success of LHCb physics program feeds on high rates and good tracking: efficiency, resolution, cleanliness – and this will not change
  - Experience has shown we cannot compromise much of what we have now in terms of performance
- Keeping the performance at FU requires solving several issues:

- Tracker (re)-structure
- Track pattern recognition
- Primary Vertex
- Physics reconstruction

- Most of this likely require new methodology, not just "more of the same"

Pattern Recognition and Tracker design

- Increasing the track density will quickly bring us over the edge where confusion is intrinsically unmanageable
  - Both ghost rate and the processing time will explode
  - Hardest issues are forward tracking, and PV in VELO

→ R*e-design tracking detectors* (costly!)

- More granularity (3D in forward?)
- Pattern-recognition-friendly geometry
- Design for ***local*** reconstruction to allow embedded reconstruction
  - e.g. CMS's double-layer design: *locally* reduces multiplicity

*A real game changer would be **time-tagged hits***

("4D tracking", ERC-winner Cartiglia docet)

How time-tagging helps PR

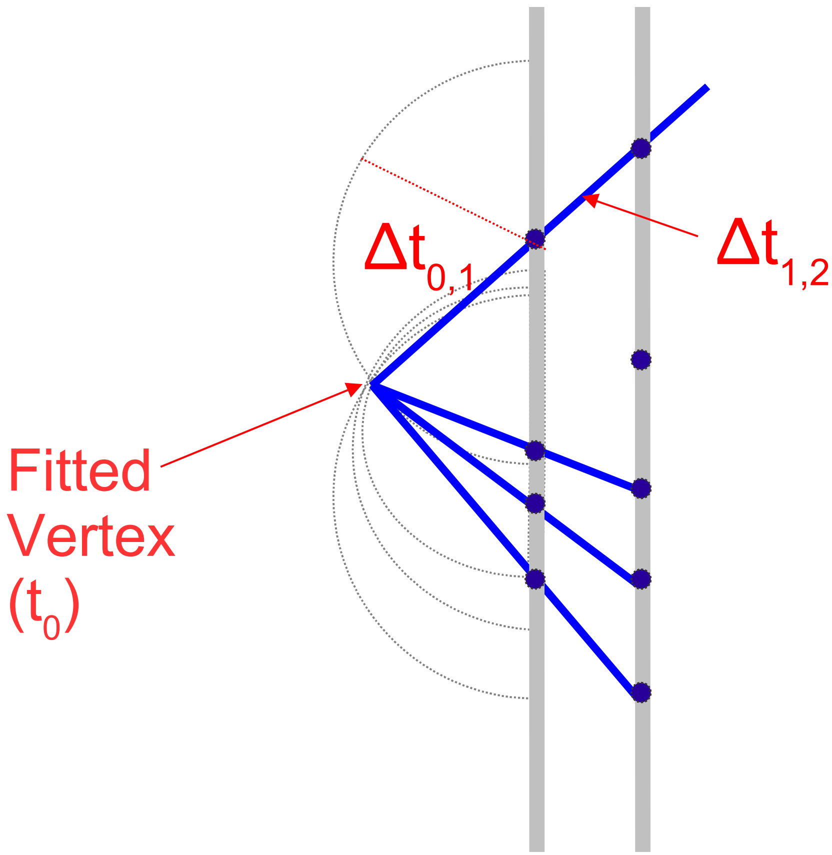

A lot of simulation work needed to evaluate the potential of this new methodology

- Timing constraint allow "<u>vertex reconstruction from a single layer</u>"
- Also strict constraints to hit association between layers (on top of bi-layer)
- May turn out to be more important that simply associating a time to a track
- R&D already exists, aimed at σ <20ps. <u>Beware</u>: costly, and not for today. Most current thoughts about a *single* timing layer.

Processing Power

- Core issue for FU tracking

If we had 40MHz readout *today* at point 8 ...

… we still could not take data at > 1MHz

That's why the online farm is such a big part of the upgrade project

- In the FU O($10^{34/35}$) we will need x40x25(0)= 10,000x the computing power of today - *and we are not allowed much time*

Said differently:

High-Lum is 10x to 100x the track density of the upgrade *at fixed 25ns*

→ >>10x to 100x the cost of the upgrade Farm in < 5 years

- Processing time **not** linear with occupancy...
- Problem compounded by the need for Real-Time analysis
- Moore's factor (2015→Phase-2 UPG) : 2^5=32 → 300x today's cost (take 2x away for "impact of hardware choices" [LHCb-INT-2013-35])

To be helped by Moore, we would need to wait until the end of the LHC

Computing performance evolution studies

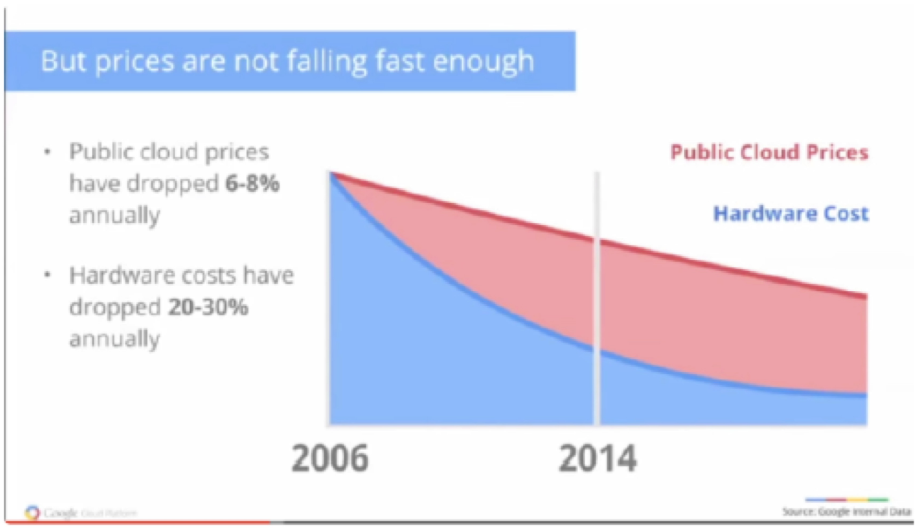

Slide from Urs Hölzle's keynote at Google Cloud Live, March 25, 2014

AWS cloud price reduction

Year	Price	Reduction	Comment
		20%	
2008	$0.800	$0.640	
2009	$0.640	$0.512	
2010	$0.512	$0.410	
2011	$0.410	$0.328	3 years, 50% reduction
2012	$0.328	$0.262	
2013	$0.262	$0.210	
2014	$0.210		3 years, 50% reduction from 2011
April 1, 2014	$0.210		6 years, 75% reduction from 2008

http://www.intel.com/assets/pdf/general/servertrendsreleasecomplete-v25.pdf

What can beat Moore's law ?

→ Improve efficiency of using Si gates

→ *Specialized tracking processors, fine-tuned to what we need*

→ *Detector-embedded tracking*

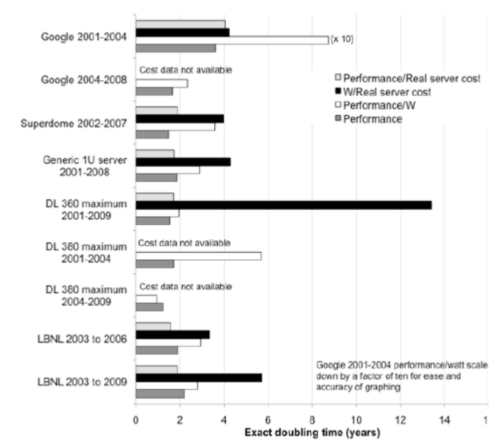

Figure ES-2: Summary of trends for servers, expressed as doubling time in ye

Longer bars mean slower growth. Doubling time calculated using instant exponential growth rates as described in the text.

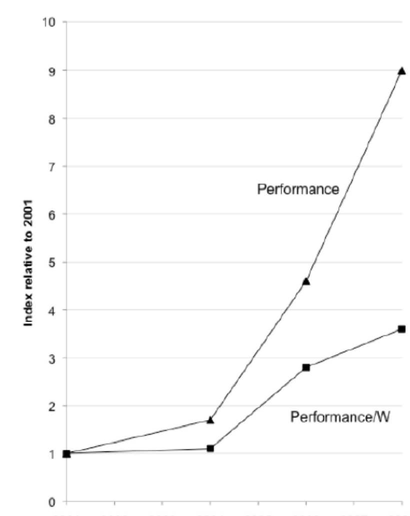

Figure 5: Performance and performance per watt trends for Google servers, 2001 to 2008

How much could we possibly gain ?

- Historical example: the Silicon Vertex Tracker of CDF (year 2000)
- Similar size and cost of the HLT farm of the time (~250 commercial CPUs)
- 30kHz vs 30 Hz: factor x1000 using the ~same silicon technology (actually SVT a bit worse due to longer TTM of fully-custom chip).
- Other examples from recent history.
    - GPU vs CPU in graphics
    - FPGA vs CPU in high-frequency trading
- Custom processor design has the potential of making FU happen
    - Note underlying silicon gates are physically the same.
    - This is not magic: by designing the circuit carefully you can make sure you use most gates of the chip at the same time (no “dark silicon”). (Much harder to do this with commercial hardware).

Gains from custom processing in FPGAs

PMC full text: Sensors (Basel). 2013 Jul; 13(7): 9223–9247.
Published online 2013 Jul 17. doi: 10.3390/s130709223

Table 3.

Calculation time comparison.

Algorithm and Platform	Execution Time	Processing Image Resolution
LSM of Ji *et al.* [3] on FPGA	15.57 ms	1,024 × 768
Chen *et al.* [40] on FPGA	2.07–3.61ms	512 × 512
Proposed Method on FPGA	15.59 ms	1,024 × 768
Direct HT Computation on PC	(a-1) 0.93 s	1,024 × 768
	(a-2) 1.26 s	1,024 × 768
	(a-3) 1.62 s s	1,024 × 768

Speedup factors of **70 ÷ 500** commonly measured in vision, military, finance applications

Table II
COMPUTING TIME OF THE HOUGH TRANSFORM

Image	Size	# edge points	Time (FPGA)	Time (CPU)	Speed-up
Figure 1(b)	512×512	33232	135.75μs	37.10ms	273.3
Figure 8(a)	1024×1024	23293	95.27μs	27.47ms	288.3
Figure 9(a)	4096×4096	80092	326.61μs	121.64ms	372.4

Primary Vertex issues

- Vertexing issue makes a big jump when going from μ=O(1) to μ=O(many)
- We should obviously be concerned about PV finding in HL scenario

→ Possible solution: use *timing* to figure out vertex – *but not guaranteed.*

→ If we can't, we may have to do everything based just on the beamline - considered as a uniform, linear source of tracks

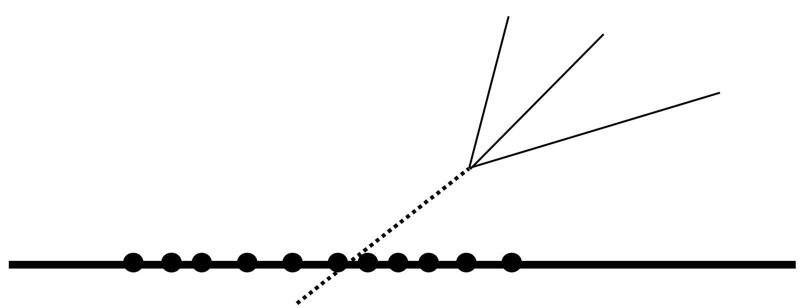

-Method: assume that the candidate momentum points back to the beamline, and take che closest point as the Z of the PV

-Simulation studies needed to evaluate what is lost in each channel

- I<u>mportant advantage</u>: freedom from the need of reconstructing anything besides the decay of interest → ***decay-oriented approach***

Conclusion

- For a Very High Lum upgrade we need to move towards:
  - *Detector-embedded reconstruction*
  - *Decay-oriented* DAQ
  - *"4D" reconstruction*
  - ***Full** analysis and calibration in real-time*
- Crucial to start immediately for the future of hadronic Flavor
- We can start this today in many ways
  - Real-time physics studies, also from RUN-2 experience
  - Simulation studies of Vertexing and 4D tracking
  - Study PR-friendly detector geometries
- Italians have been pionereeing this field
  [see "WhatNext white paper" Extreme Flavor scenario
  for details and ideas: http://www.lnf.infn.it/sis/frascatiseries/volumes.html]