ATLAS ONLINE DETERMINATION AND FEEDBACK OF LHC BEAM PARAMETERS

Josh Cogan
Emanuel Strauss
Rainer Bartoldus
David Miller

For the ATLAS Collaboration

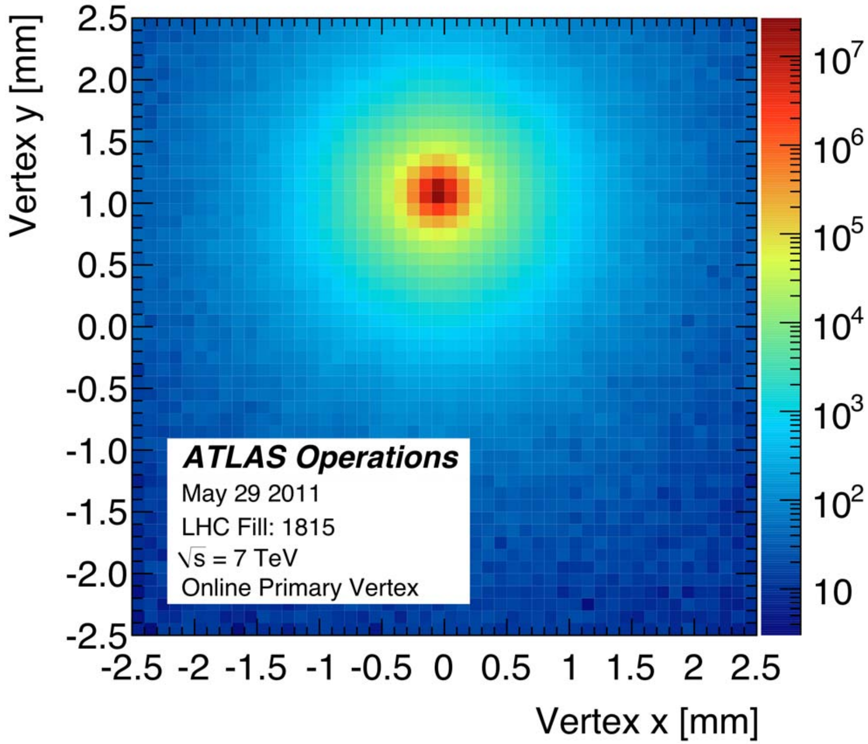

NATIONAL ACCELERATOR LABORATORY

BASIC PROBLEM 1/2

- ATLAS trigger algorithms use the beam spot to maintain higher efficiency of interesting events
  - Beam spot: location and size of luminous region

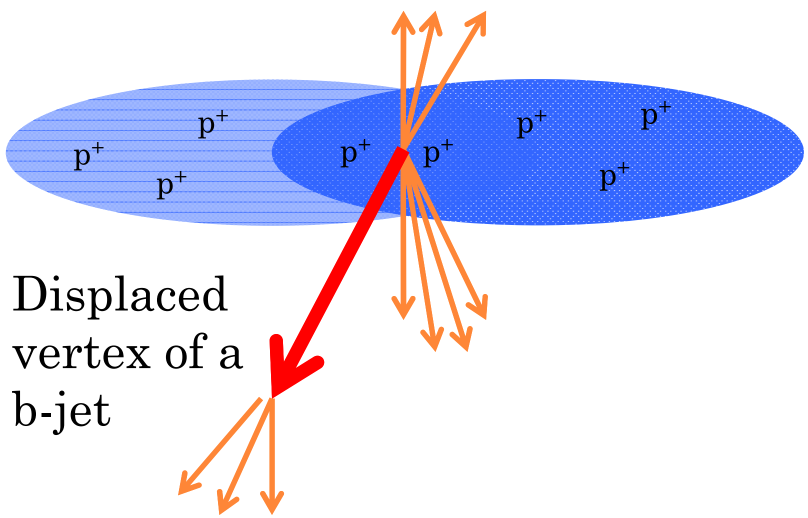

We provide with errors:
Ellipsoid Mean (xyz)
Ellipsoid Width (xyz)
Ellipsoid Tilt (xz, yz)

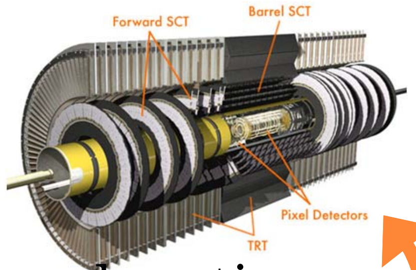

  - Used for tracking algorithms and displaced vertices
  - Measure via distribution of charged particle vertices found from hits on ATLAS silicon detectors
- Precise measurement of beam phase space at interaction point (IP)

BASIC PROBLEM 2/2

- But, luminous region changes during a fill

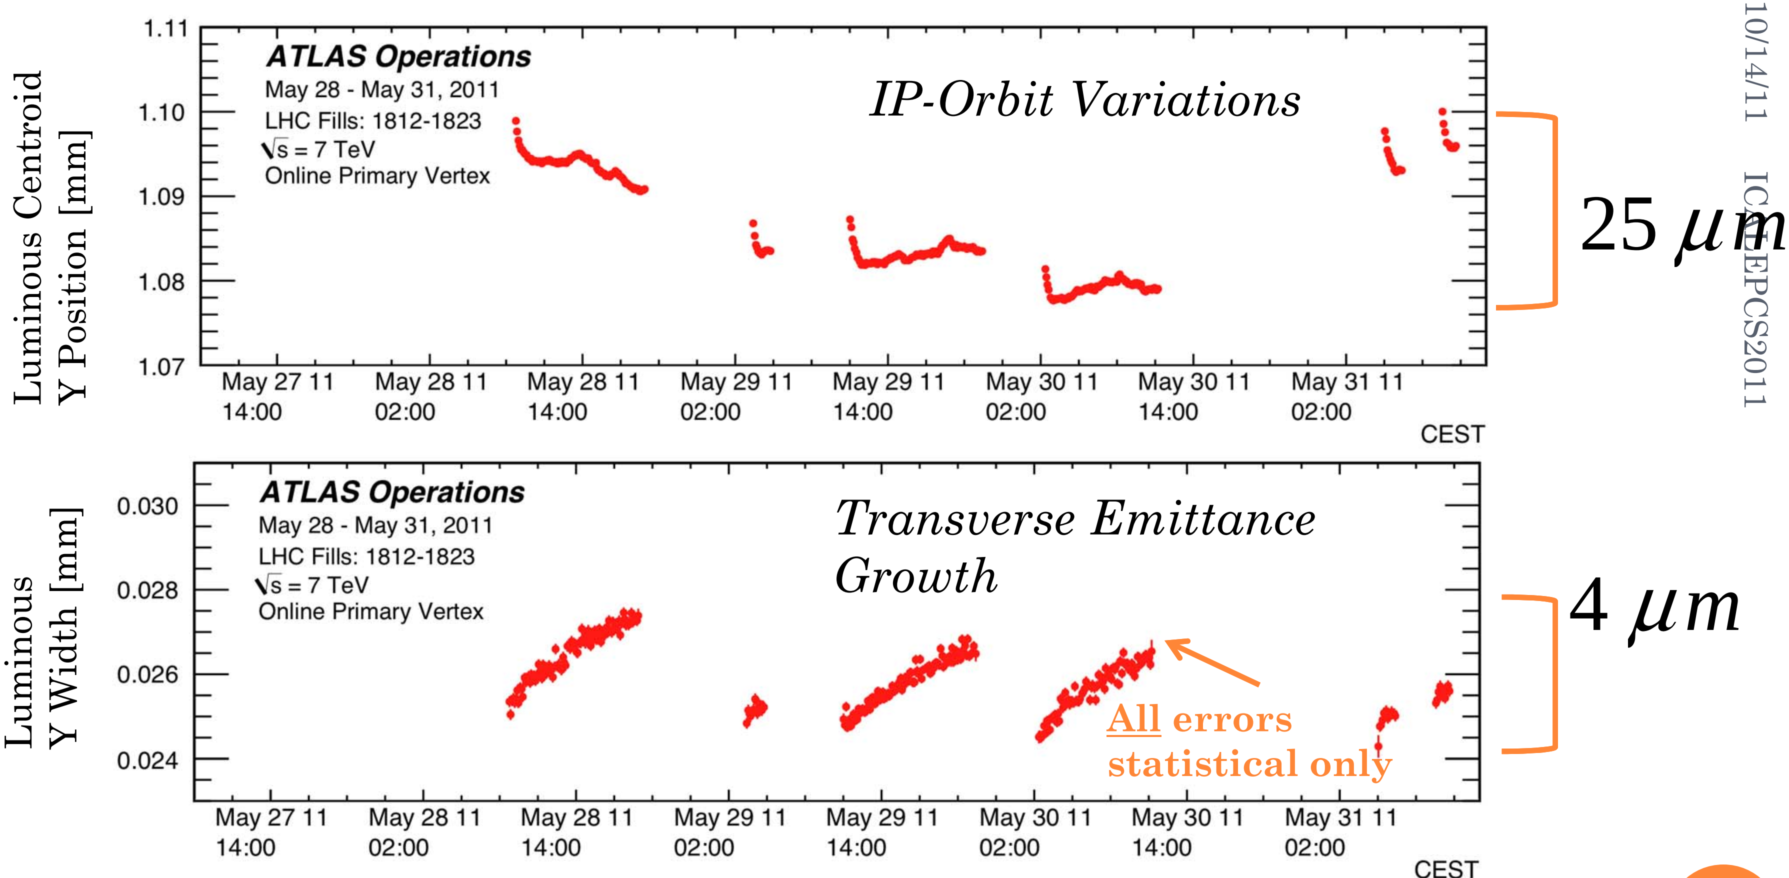

- The High Level Trigger **needs feedback**!
  - "Software" trigger working on Linux server farm

HURDLES

Focus of this talk

Algorithmic:
- Vertex resolution ~25 μm, but beam spot < 20 μm
- Operate on the trigger farm: limited bandwidth and CPU
- Only one chance to use event
- One event has many vertices!

Commissioning:
- Not in the original design
- Like changing the engine in a moving car
- Takes stable beams to test full system and feedback

Communication:
- Calculating beam spot needs **> 100,000 vertices** for 1300 bunches
- **13,000 processes** need to know beam spot
- Cannot read out entire detector at the hardware trigger rate
- Shouldn't disrupt data taking

PILE UP VERTICES

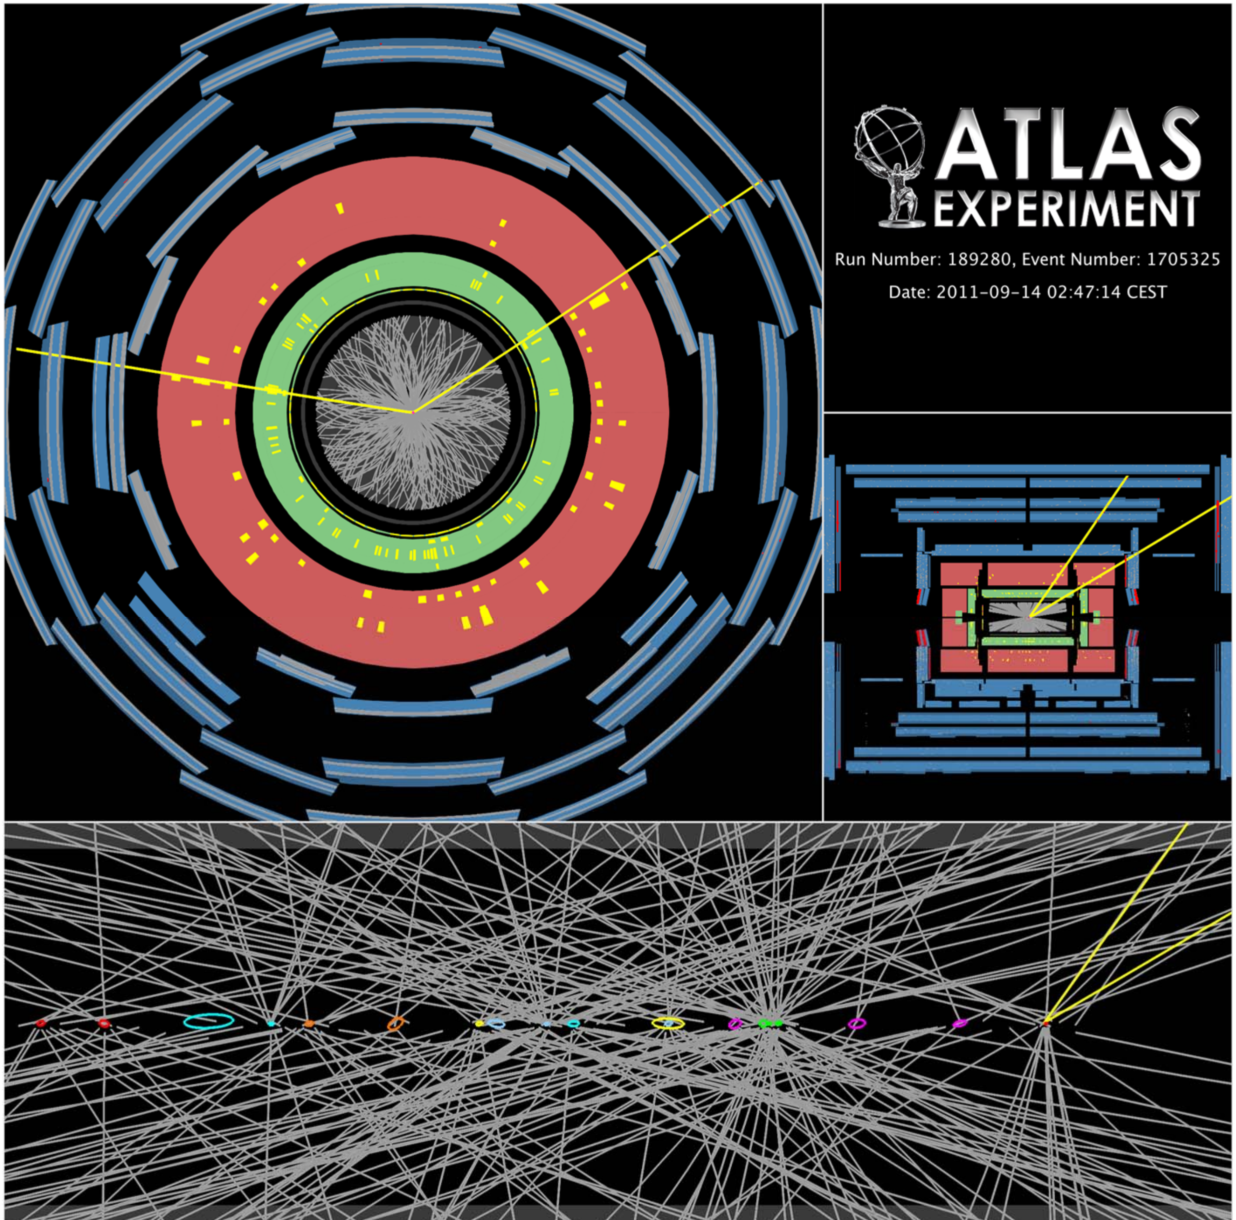

- At current luminosities there are 15-20 vertices per bunch crossing!
  - “Pile-up”
- Many vertices to fit! However…
- Computationally extremely expensive to reconstruct in real time

GENERAL SOLUTIONS

Focus of this talk

Algorithmic:
- High rate/quality of vertices
- Specialized resolution determination via "split vertex"
- Use pile up vertices as well
- Share bandwidth/CPU with other tracking intensive algorithms

Commissioning:
- Emulate online system for test and development
- LHC down time → test changes
- Special data taking calibration stream

Communication:
- Parallelize Parallelize Parallelize!
- **Fan In/Out** calculations' input and output to central locations
- Piggy back on **event data**

SOLUTION OVERVIEW

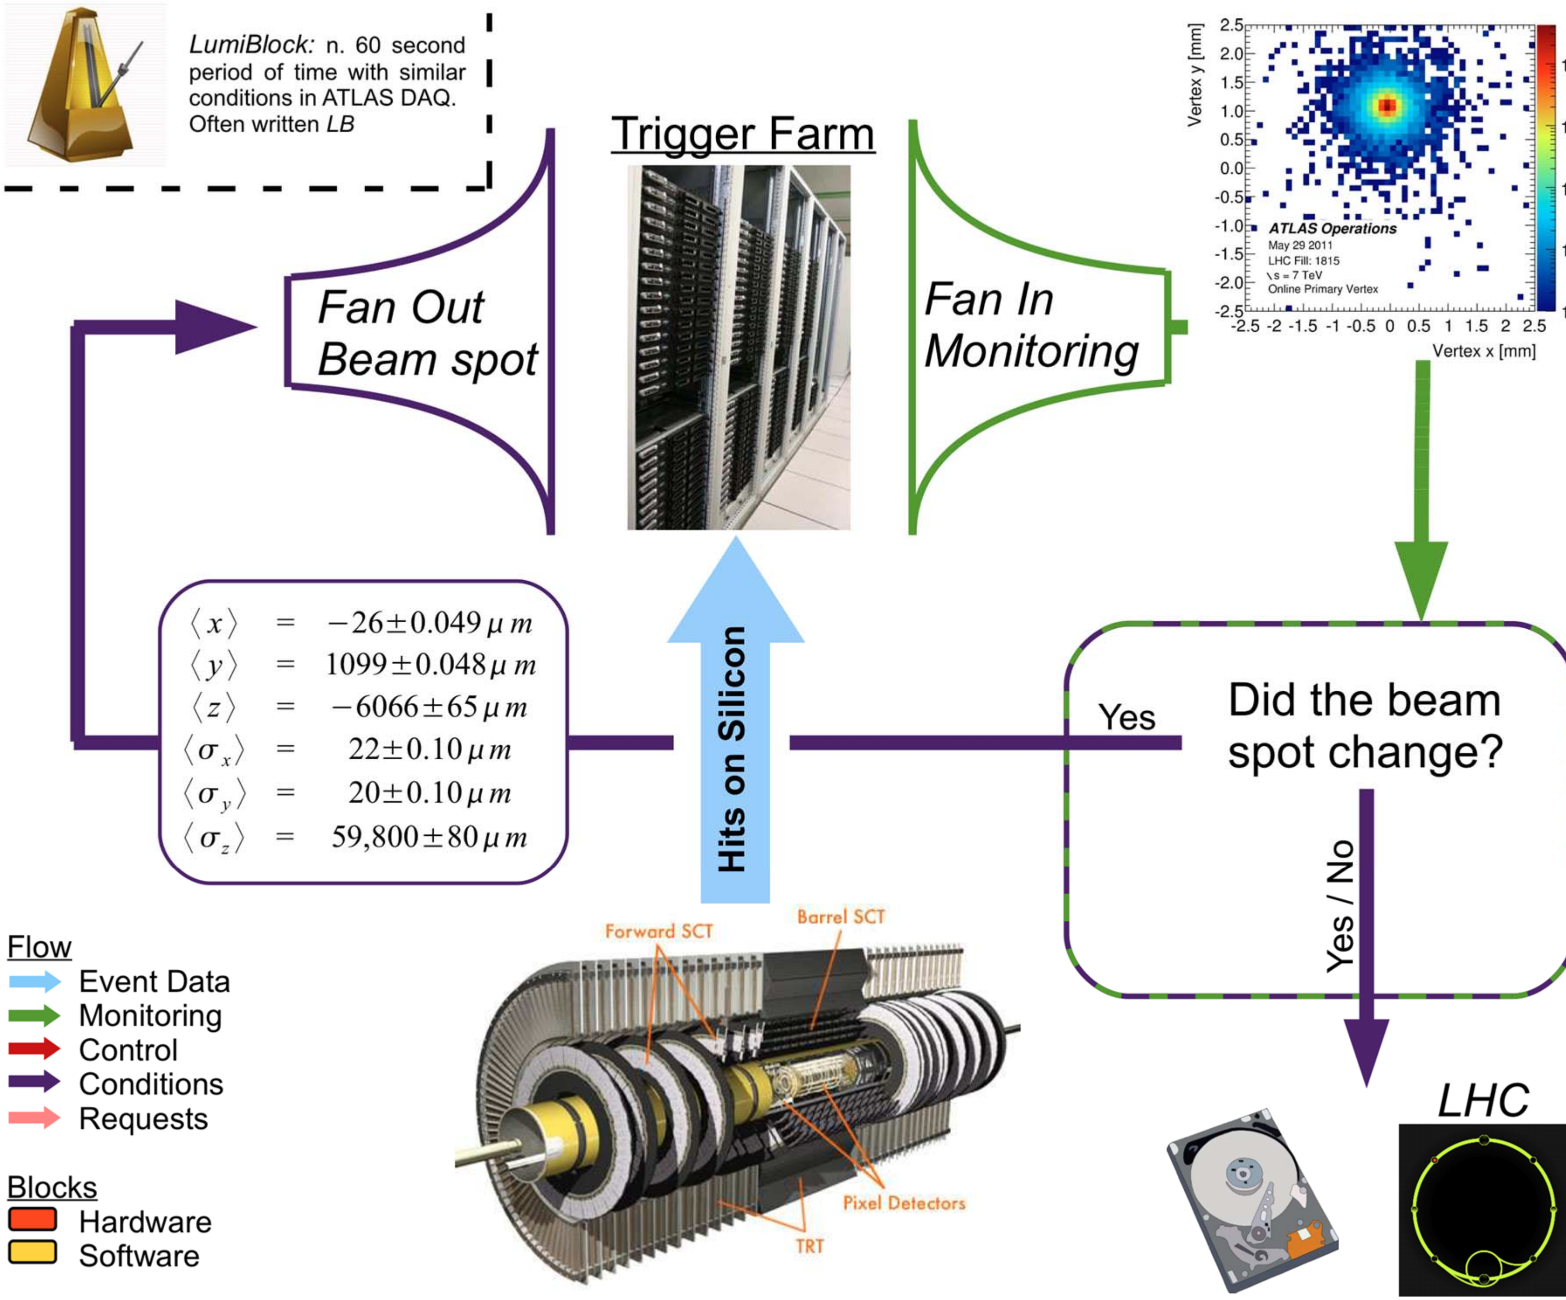

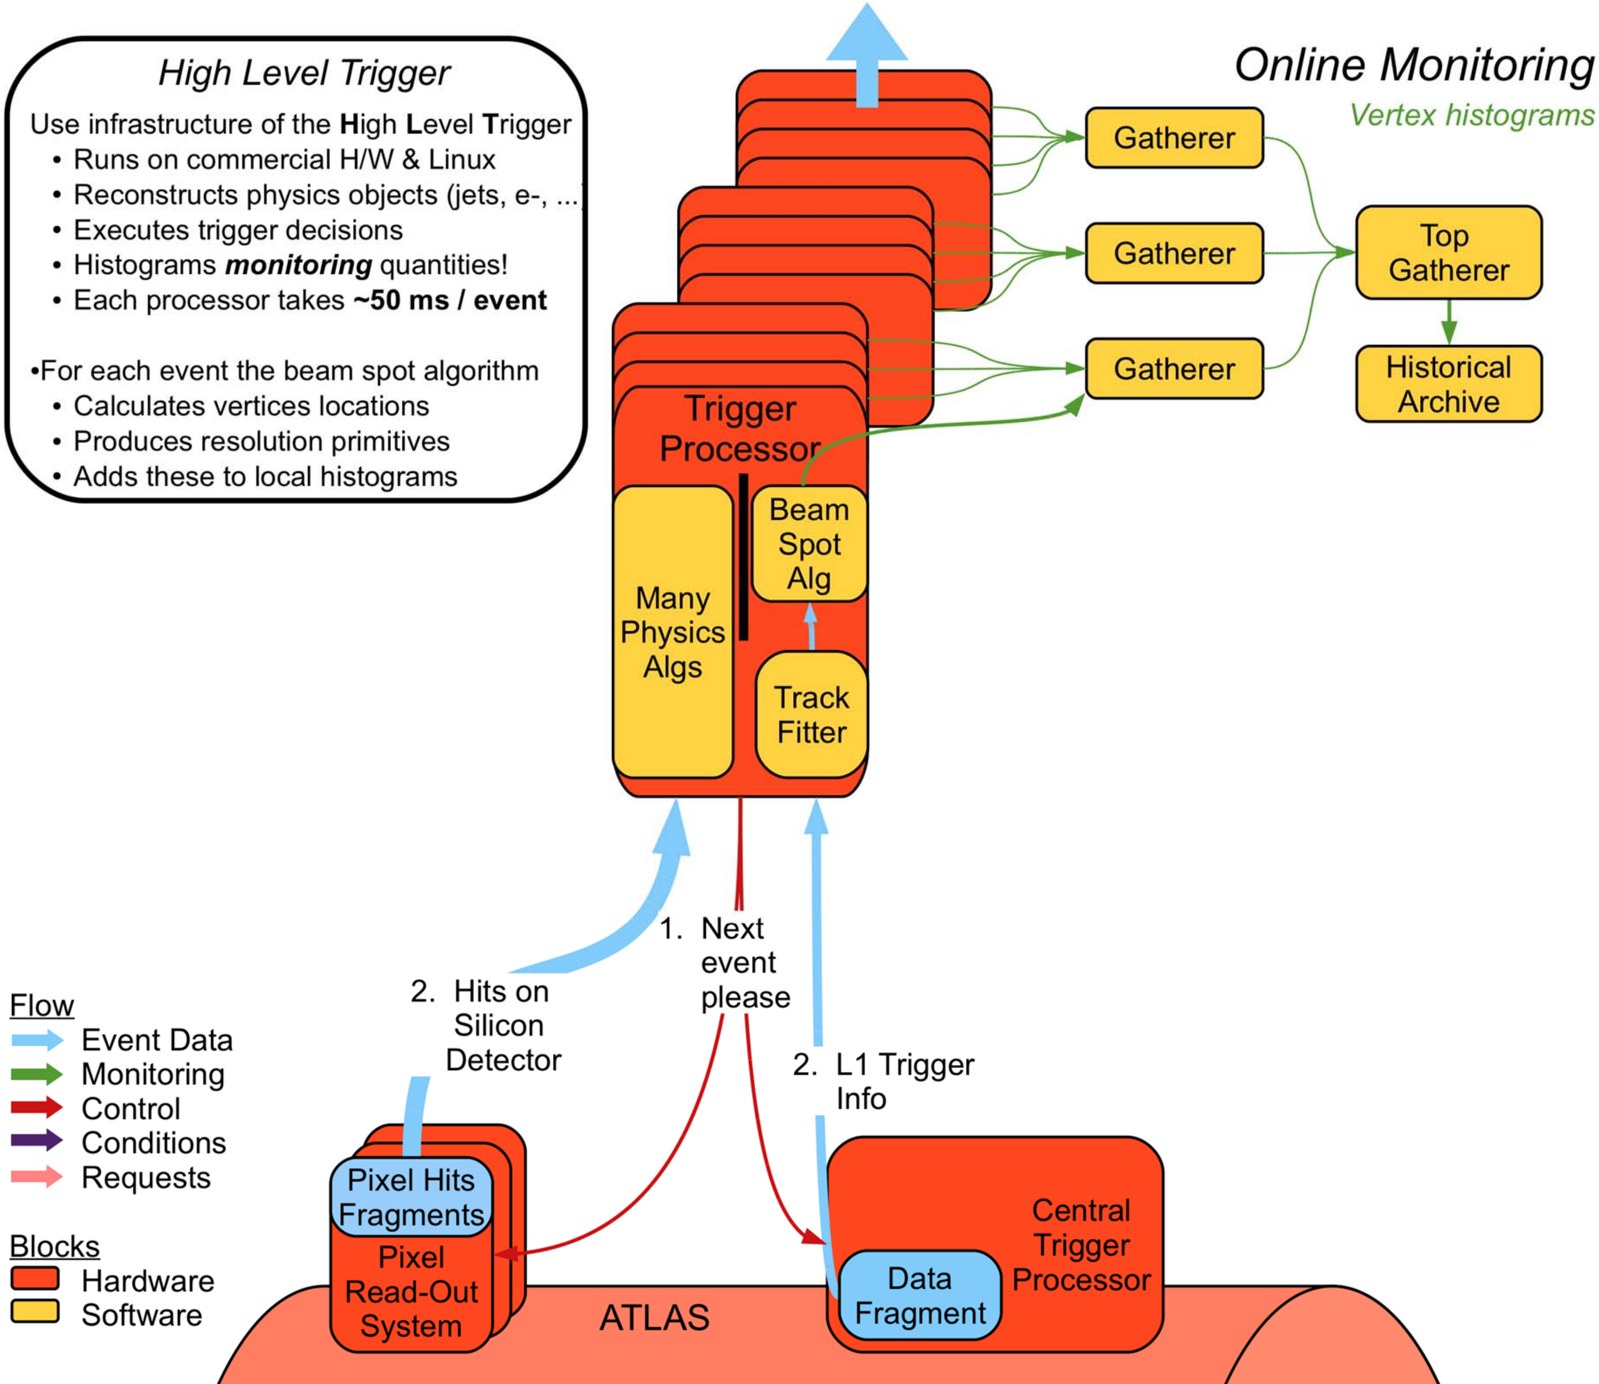

High Level Trigger
Use infrastructure of the High Level Trigger
• Runs on commercial H/W & Linux
• Reconstructs physics objects (jets, e-, ...
• Executes trigger decisions
• Histograms monitoring quantities!
• Each processor takes ~50 ms / event
•For each event the beam spot algorithm
• Calculates vertices locations
• Produces resolution primitives
• Adds these to local histograms
Online Monitoring
Vertex histograms
Gatherer
Gatherer
Gatherer
Top Gatherer
Historical Archive
Trigger Processor
Many Physics Algs
Beam Spot Alg
Track Fitter
1. Next event please
2. Hits on Silicon Detector
2. L1 Trigger Info
Pixel Hits Fragments
Pixel Read-Out System
ATLAS
Data Fragment
Central Trigger Processor
Flow
Event Data
Monitoring
Control
Conditions
Requests
Blocks
Hardware
Software

High Level Trigger

Use infrastructure of the **H**igh **L**evel **T**rigger
- Runs on commercial H/W & Linux
- Reconstructs physics objects (jets, e-, ...
- Executes trigger decisions
- Histograms ***monitoring*** quantities!
- Each processor takes **~50 ms / event**

•For each event the beam spot algorithm
- Calculates vertices locations
- Produces resolution primitives
- Adds these to local histograms

Gatherer

Processes needs to cooperate!
- Need O(10^5) vertices for a beam spot
- Each event/processor are independent
- Merge those histograms across farm

Over 6,000 processors: how to merge?
- Aggregate at rack and farm level
- Rack: merge ~240 processors
- Top: merge ~30 racks

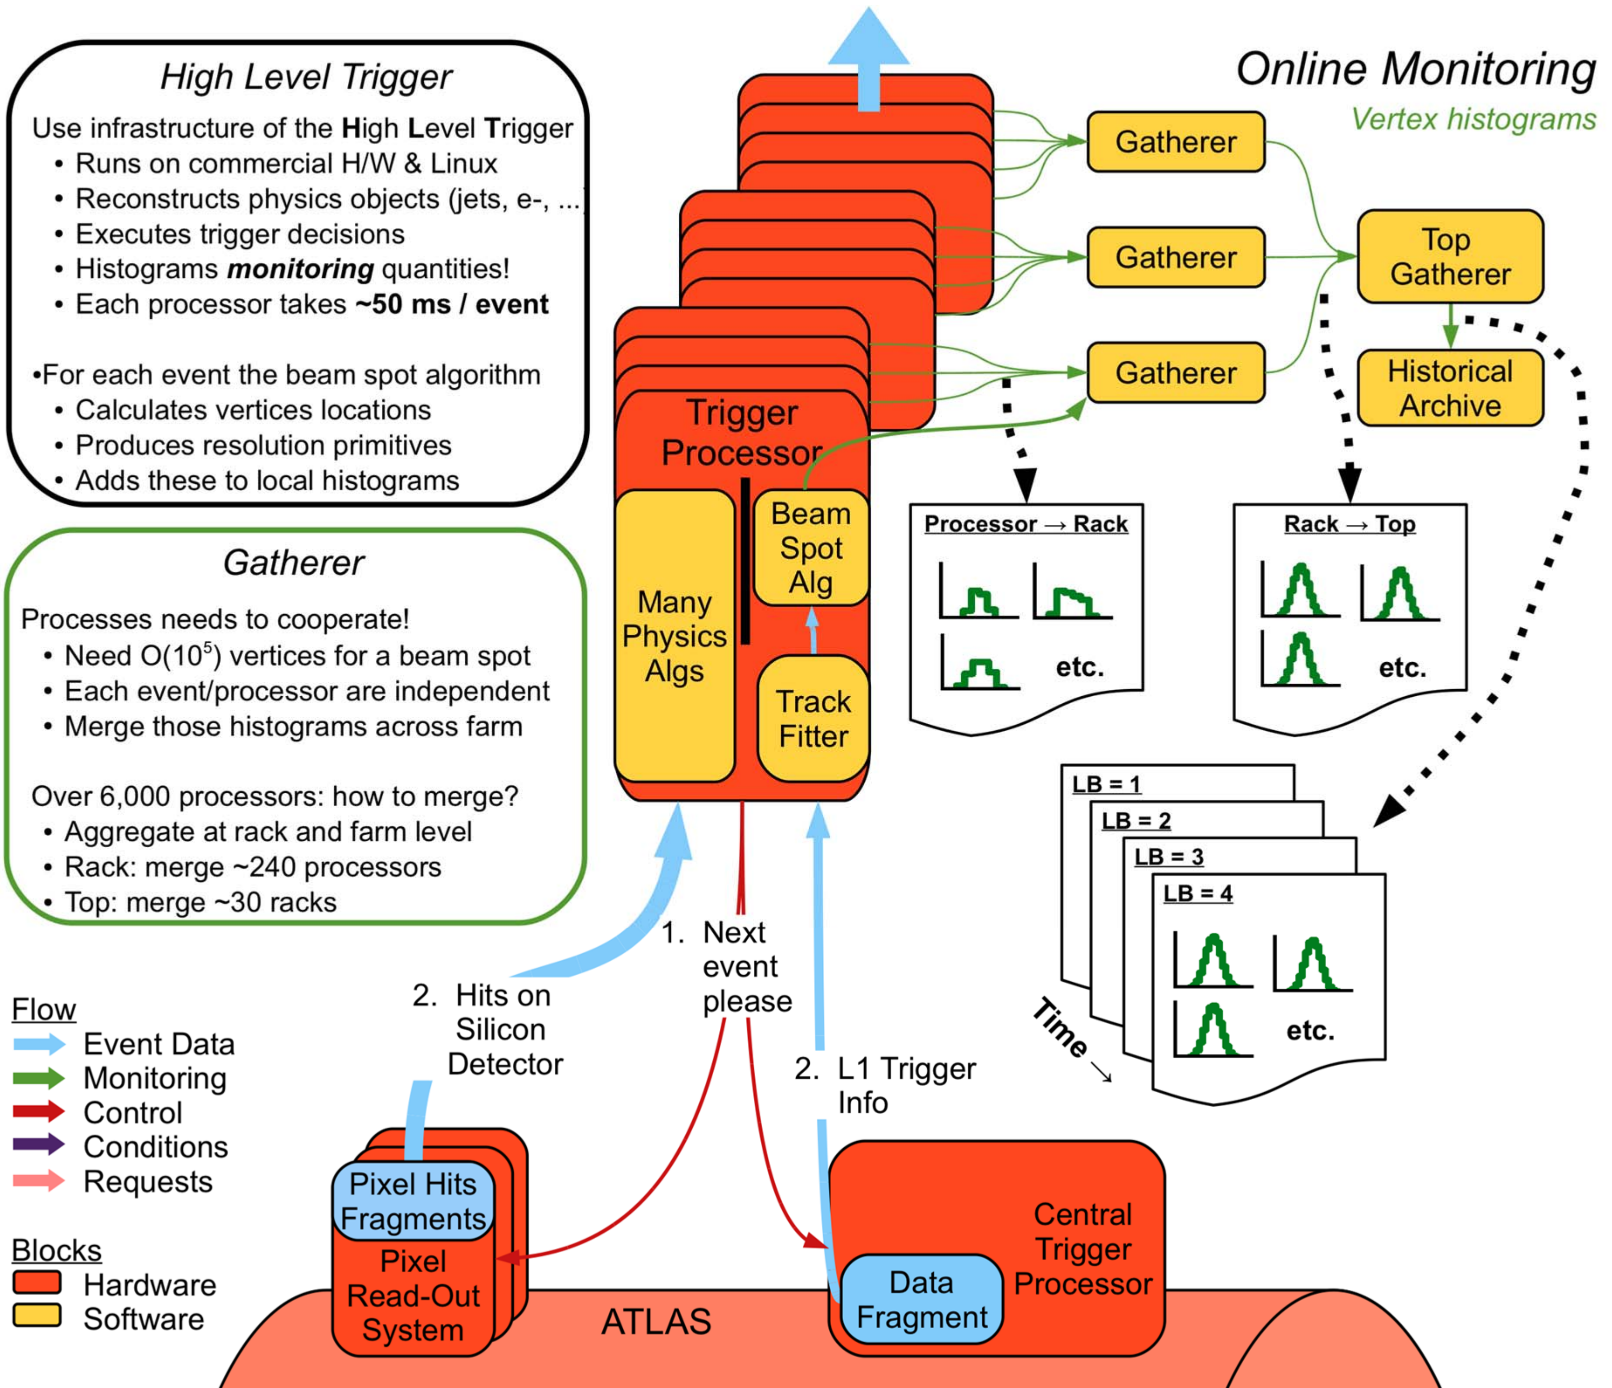

Online Monitoring

Vertex histograms

Calculation

Histograms → Beam spot
- Input data now centralized
- Calculate beam spot from histograms
  - Fit gaussians, calculate resolutions, ...
- Write values to file and send to LHC
- Is there a significant difference between current and nominal values?

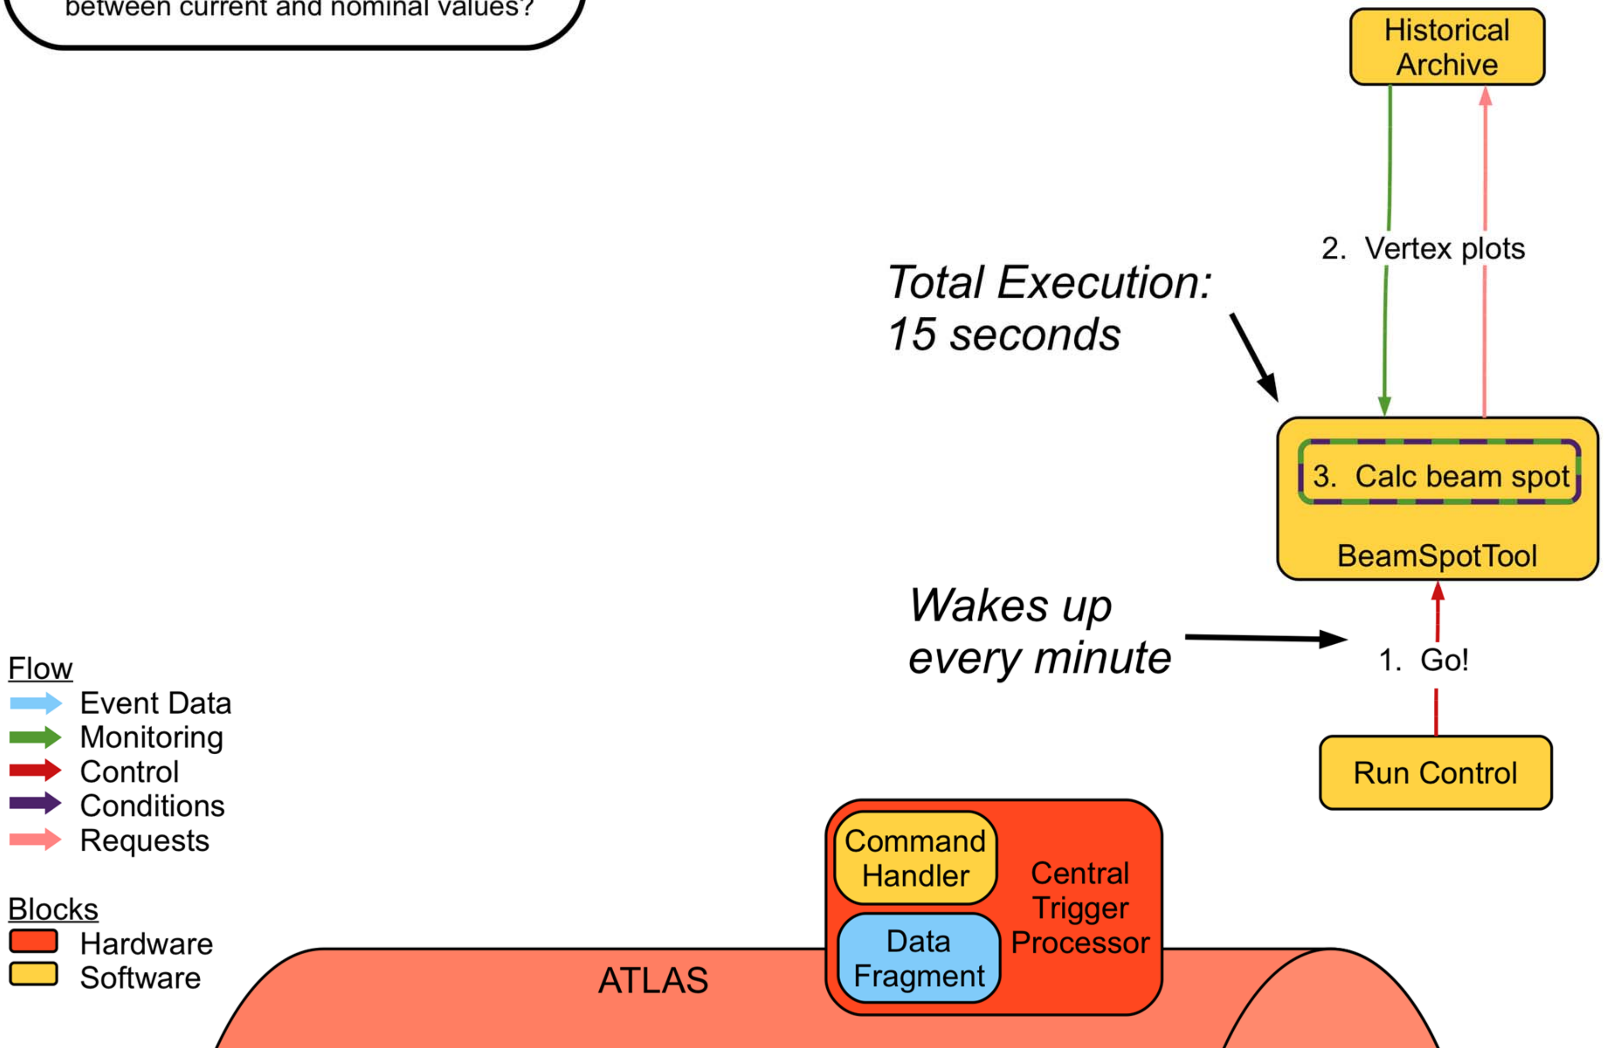

Online Monitoring

Vertex histograms

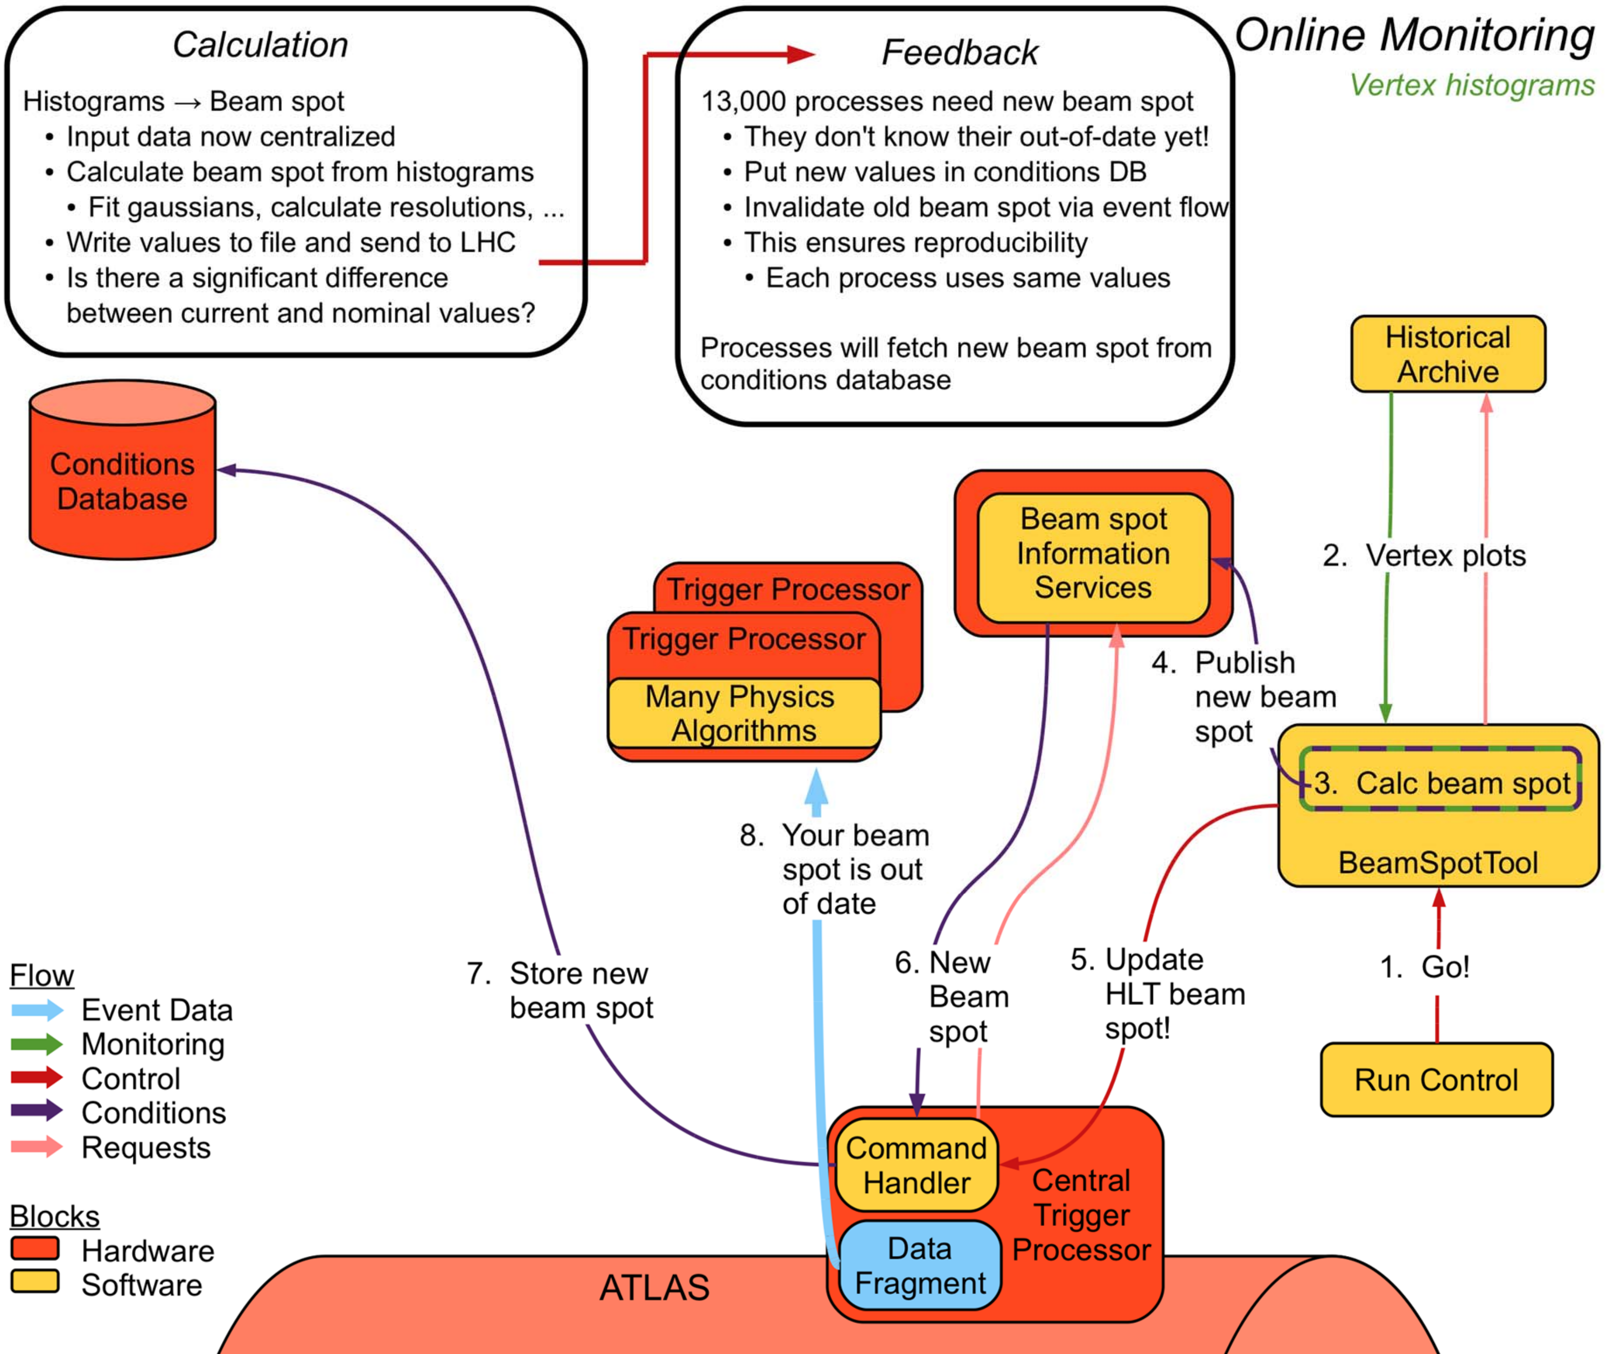

FEEDBACK CRITERIA

- Compare two sets of beam spot parameters
  - *Current*: from histograms just out of **trigger farm**
  - *Nominal*: from the last update--stored in conditions DB and used by the trigger farm for tracking algorithms
- Decide to update (feedback) if:
  1. Position offset > 10% width
  2. Width offset > 10% of itself
  3. Error on any measurement decreases by 50%
  4. Nominal is invalid (and current is valid) ← **We invalidate after beam dump**
- Criteria are completely configurable!
  - Meet the needs of clients but easy to do better

Conditions Database

Beam spot parameters

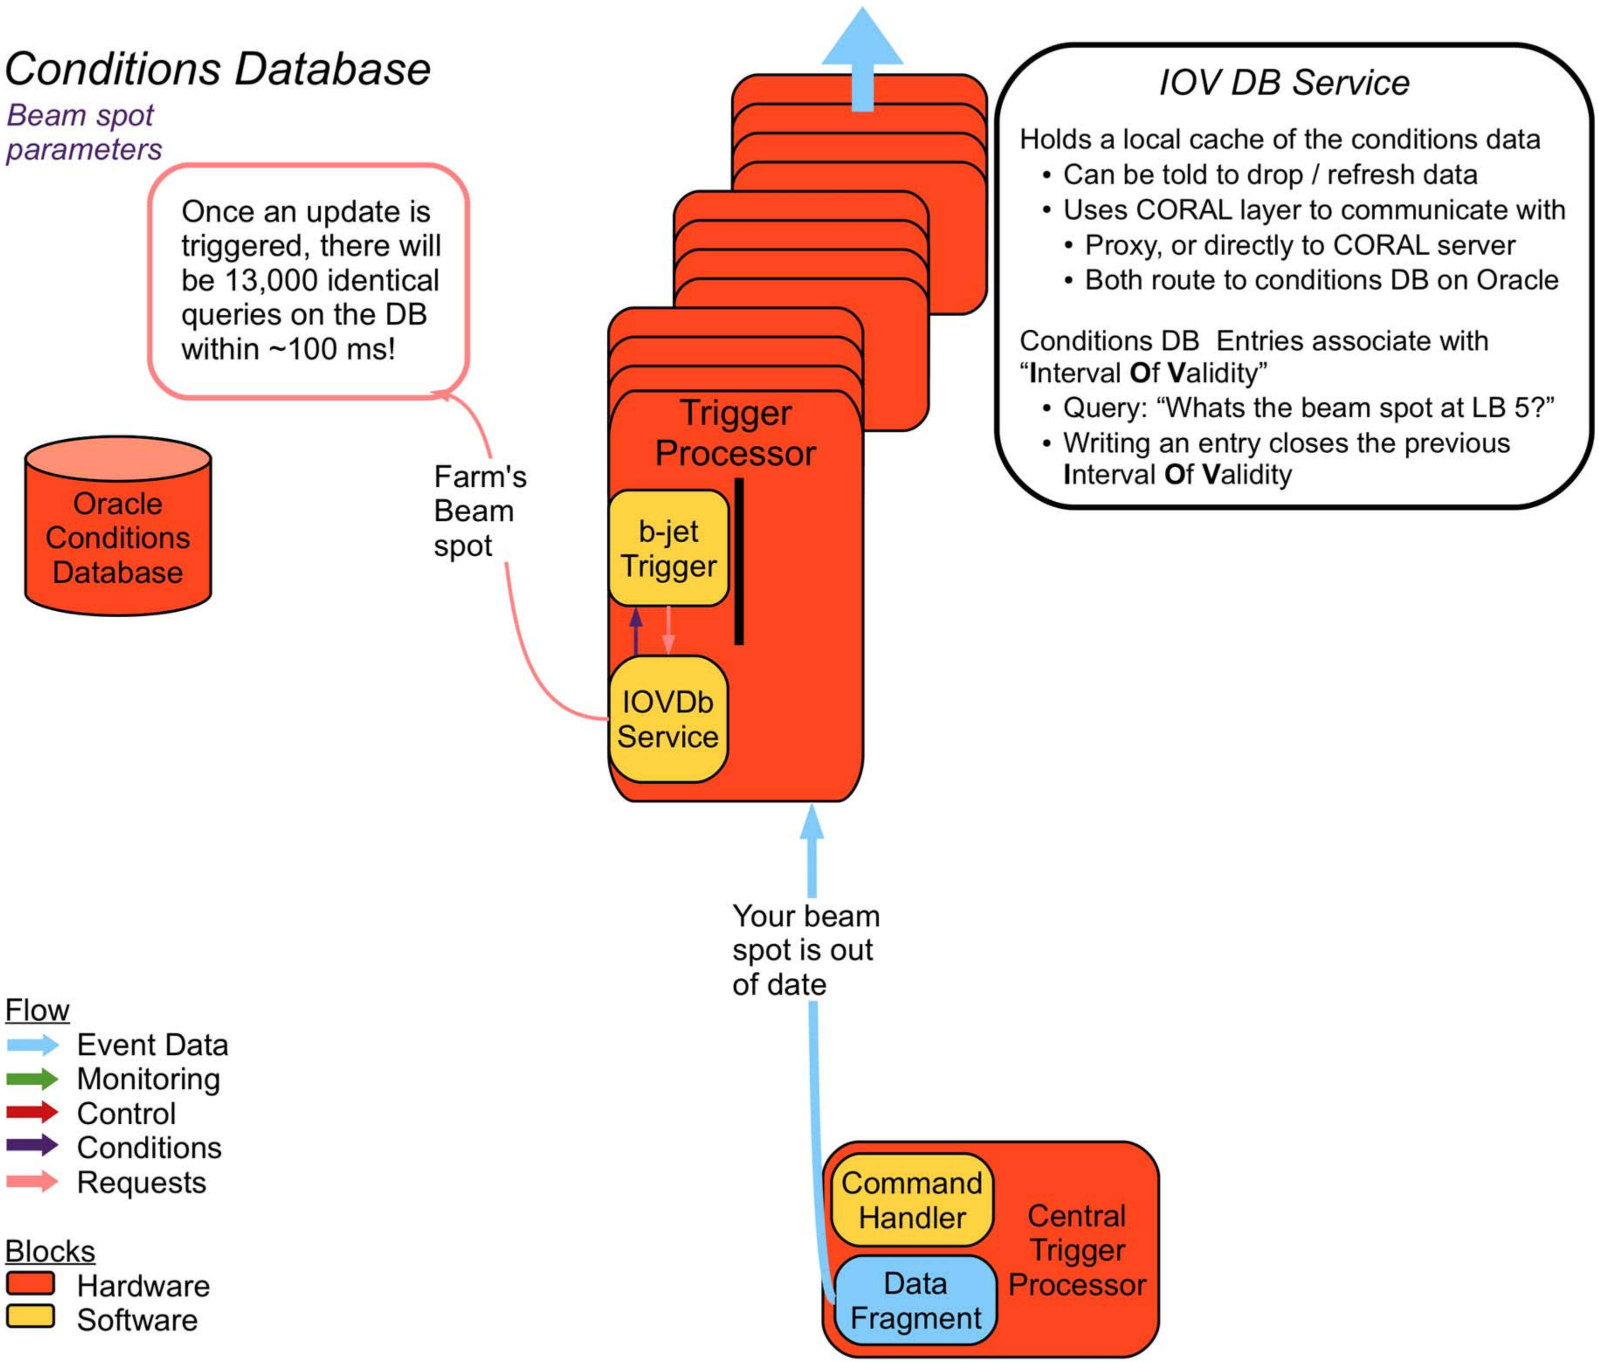

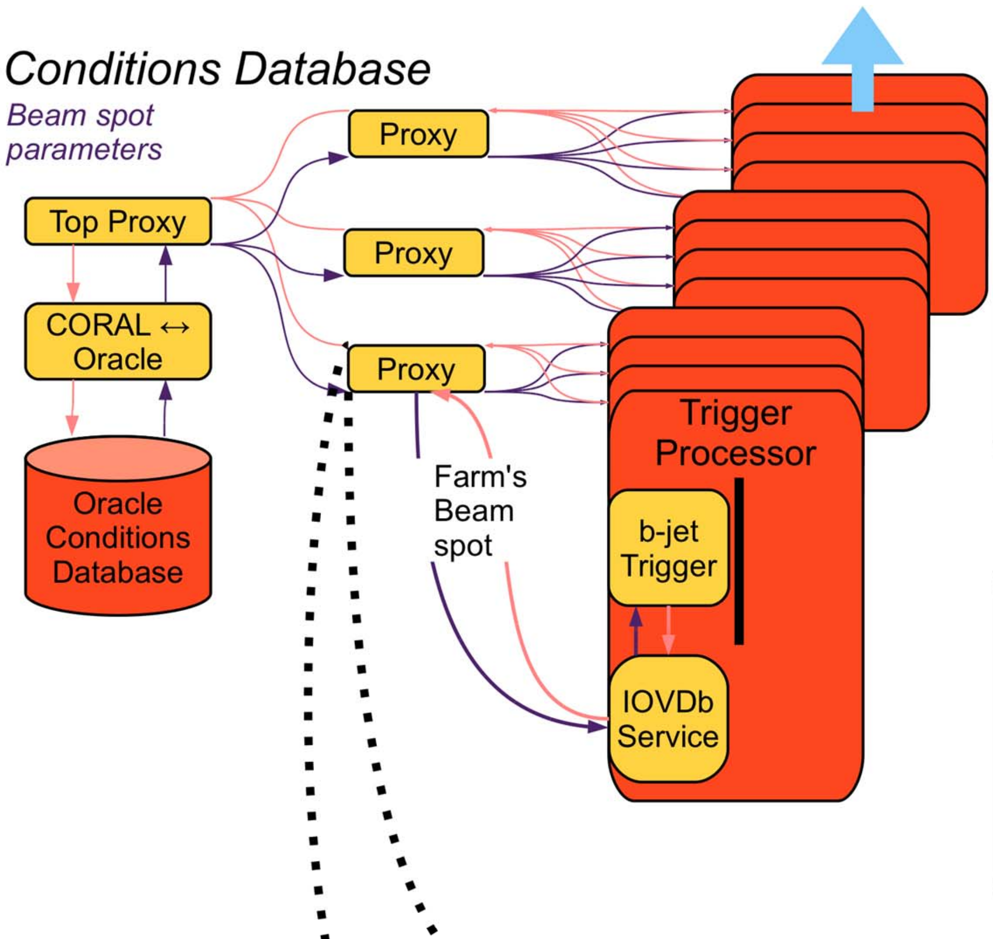

IOV DB Service

Holds a local cache of the conditions data

- Can be told to drop / refresh data
- Uses CORAL layer to communicate with
  - Proxy, or directly to CORAL server
  - Both route to conditions DB on Oracle

Conditions DB Entries associate with "**I**nterval **O**f **V**alidity"

- Query: "Whats the beam spot at LB 5?"
- Writing an entry closes the previous **I**nterval **O**f **V**alidity

CORAL Proxies

Nodes ask same question at same time

- Just like gathering, use a multiplex tree
- Each proxy looks like a server & client
- Whole tree capable of answering 10000's of identical request in ~10 ms

CORAL is a open server/client API

- High performance routing and caching
- Decouples HLT from DB engine

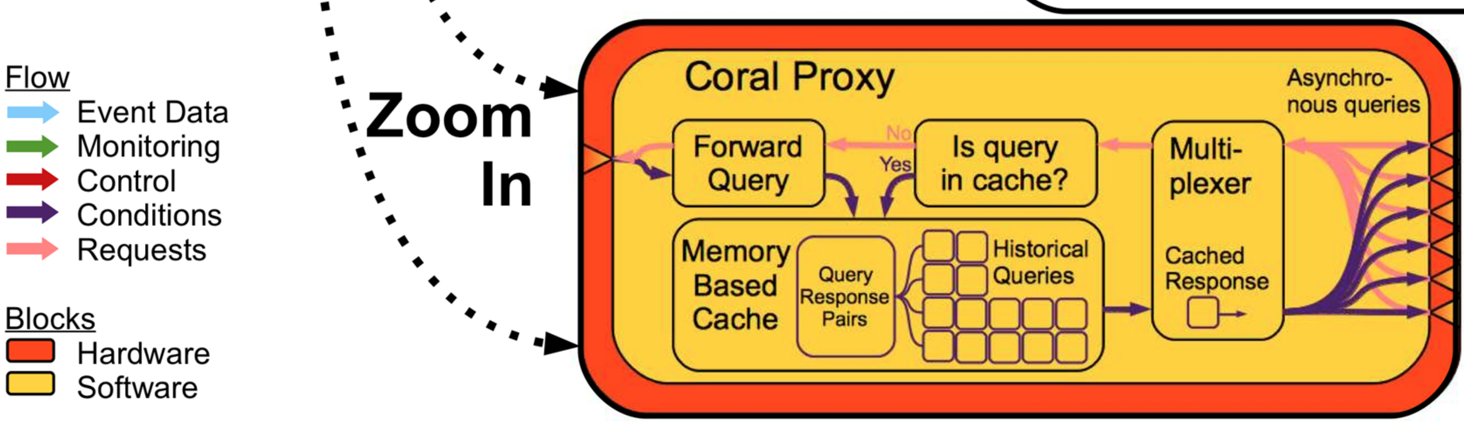

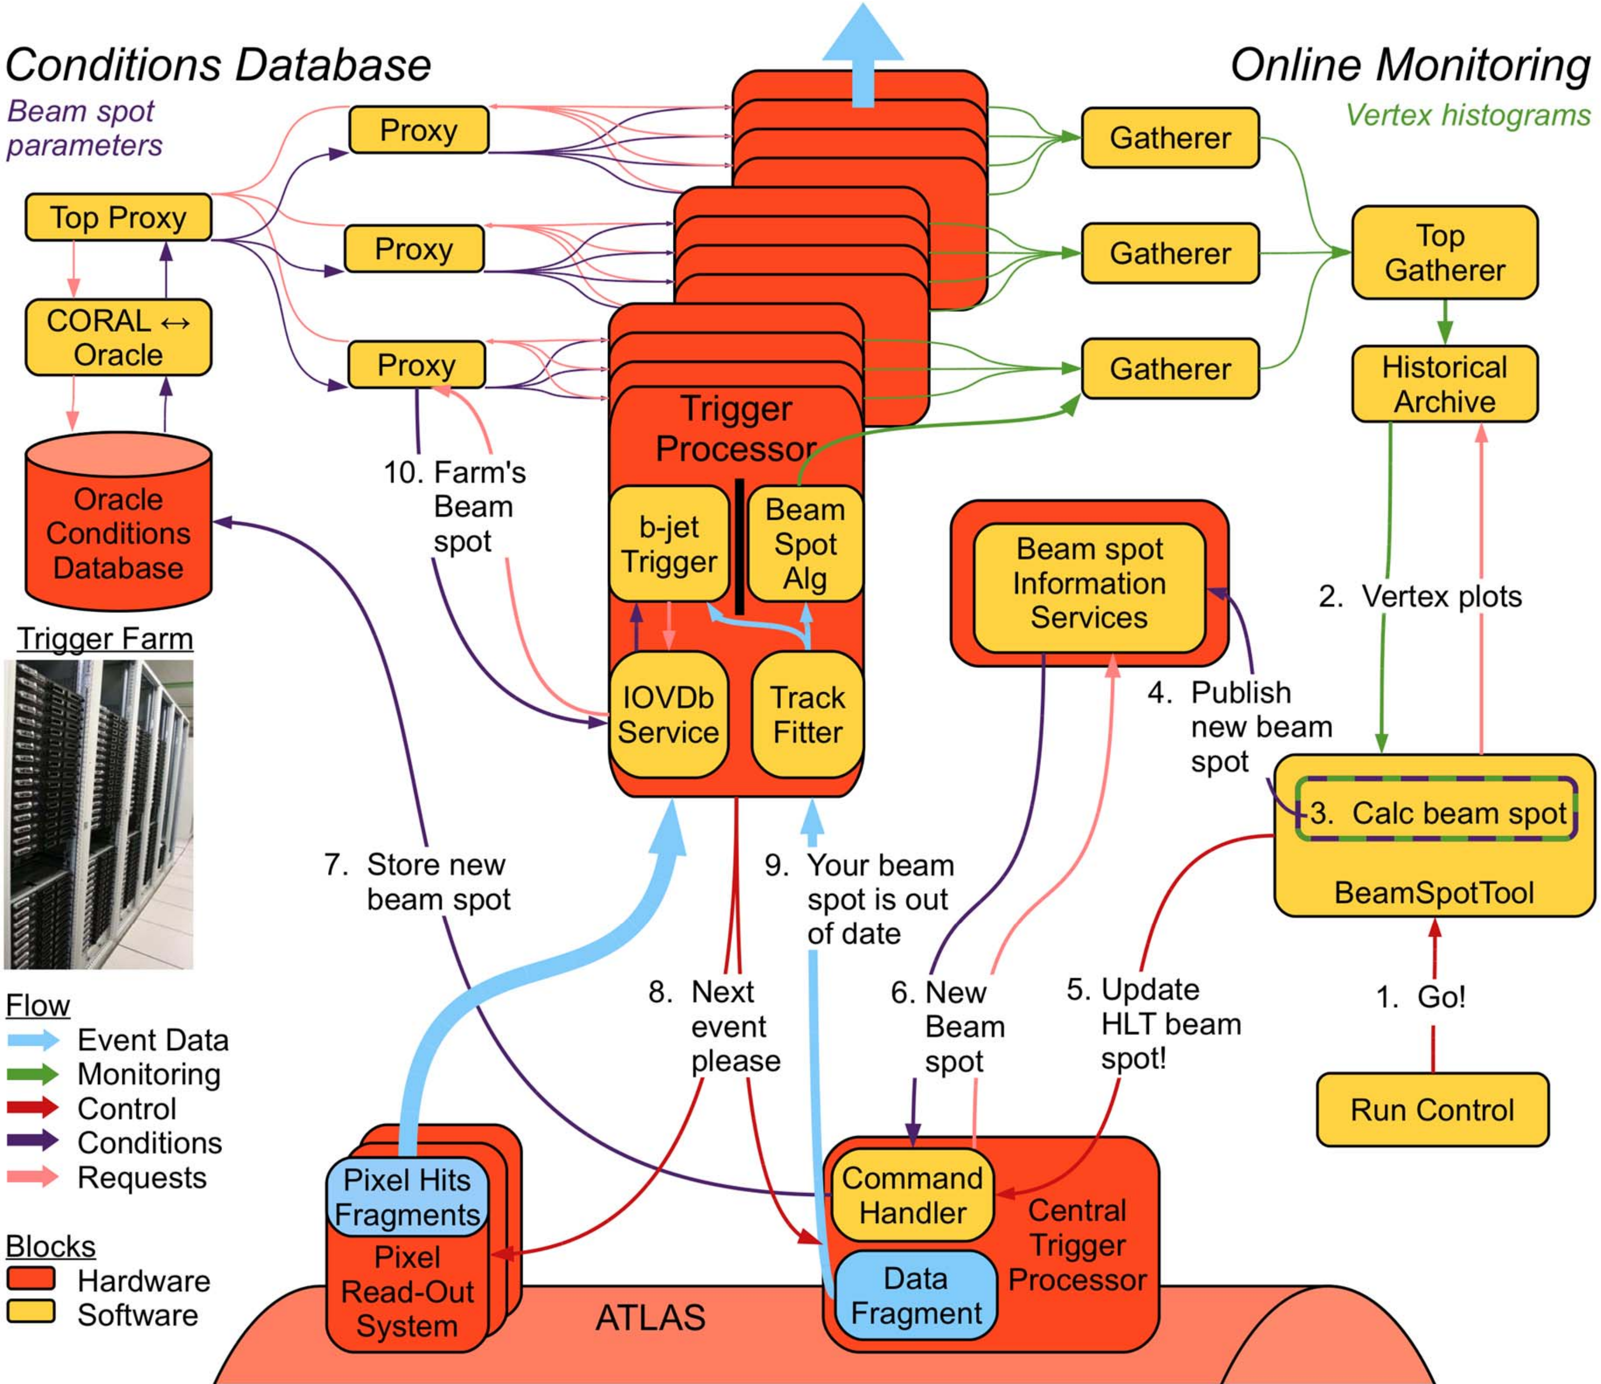
Conditions Database
Beam spot parameters
Online Monitoring
Vertex histograms
Proxy
Proxy
Proxy
Top Proxy
CORAL ↔ Oracle
Oracle Conditions Database
Gatherer
Gatherer
Gatherer
Top Gatherer
Historical Archive
Trigger Processor
b-jet Trigger
Beam Spot Alg
IOVDb Service
Track Fitter
Beam spot Information Services
10. Farm's Beam spot
2. Vertex plots
4. Publish new beam spot
3. Calc beam spot
BeamSpotTool
Trigger Farm
7. Store new beam spot
9. Your beam spot is out of date
8. Next event please
6. New Beam spot
5. Update HLT beam spot!
1. Go!
Run Control
Flow
Event Data
Monitoring
Control
Conditions
Requests
Blocks
Hardware
Software
Pixel Hits Fragments
Pixel Read-Out System
ATLAS
Command Handler
Data Fragment
Central Trigger Processor

RESULTS: FEEDBACK 1/2

- Latency ~ 240 seconds: Analyzing event→feedback
  - 140 seconds: Gathering at fixed intervals
  - 60 seconds: Waiting for update at LB increase
  - **15 seconds**: Fitting Gaussians, calculating beam spot
  - Could force updates but 5 minutes fast compared to beam
- Frequency of actual feedback. N.B. *not* periodic!
  - At start of run need to bootstrap (start from scratch)
    - Prevents large tracking errors in case beam moved significantly
    - Errors on the values drop rapidly as statistics grow
    - First update 5 minutes after data taking starts (*invalid before this*)
    - ~4 updates in first 25 minutes of data taking
  - During the fill, beam changes slowly
    - Emittance blow up, IP orbit variations …
    - ~1 update every few hours after bootstrap phase

RESULTS FEEDBACK: 2/2

- Process pause: ~10 ms to fetch new beam spot
  - Proxy tree & event time stagger → most don't wait
  - No deadtime! No DAQ busy!
- Difference between *current* and *nominal*

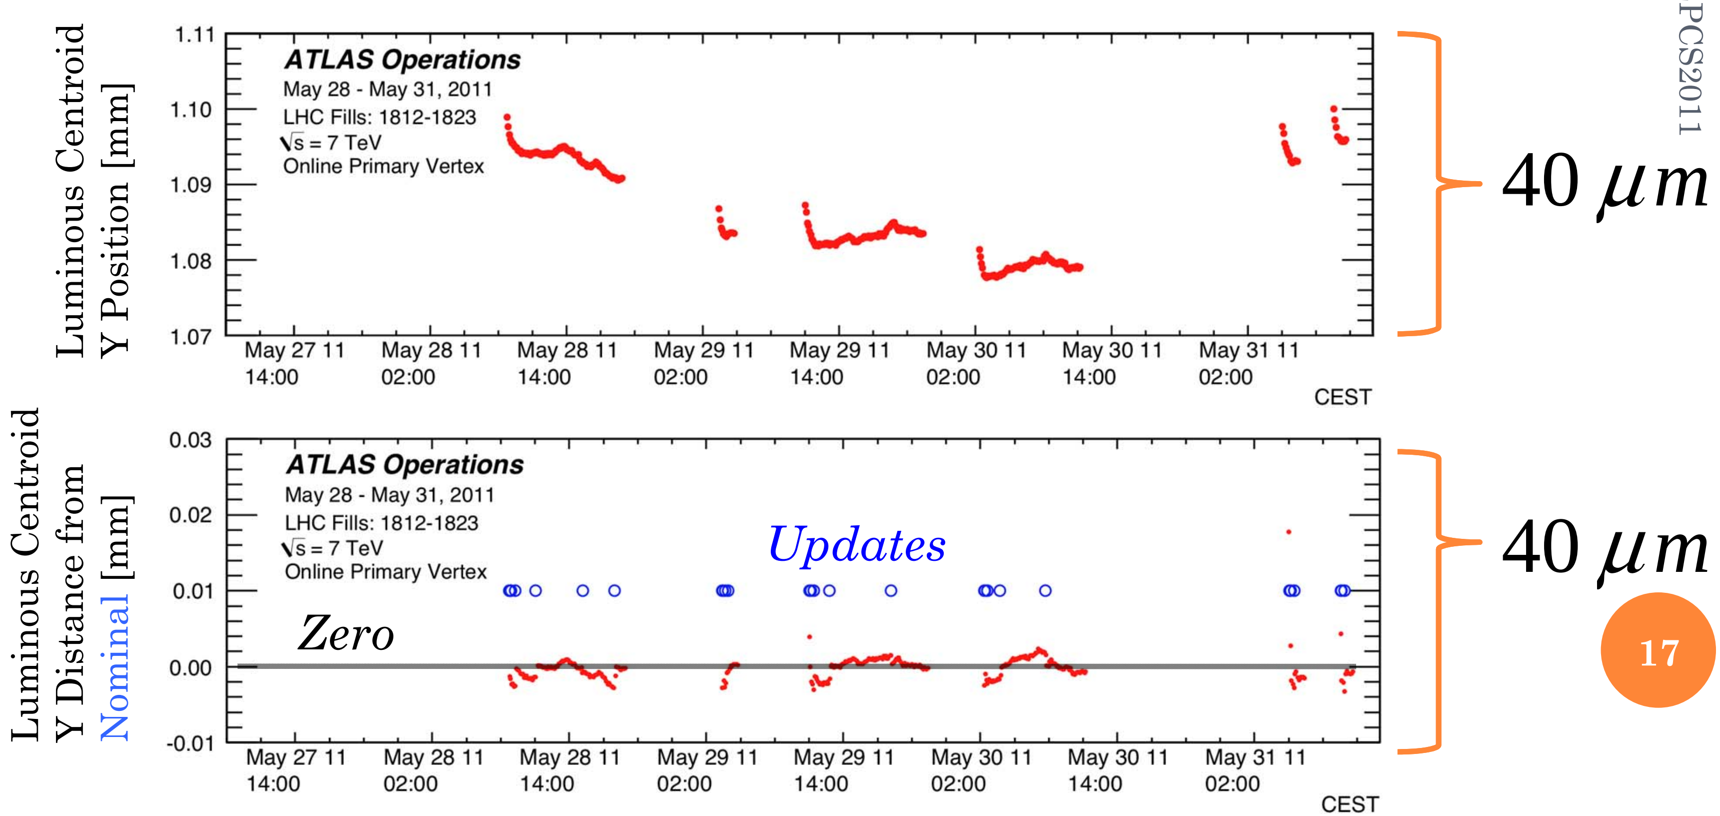

LHC CONFIGURATION PAGE

04-Oct-2011 21:14:23 Fill #: 2182 Energy: 3500 GeV I(B1): 1.76e+14 I(B2): 1.77e+14

Accelerator Mode:	PROTON PHYSICS	Beam Mode:	STABLE BEAMS

Active Filling Scheme:	50ns_1380b+1small_1318_39_1296_144bpi
Active Hypercycle:	3.5TeV_10Aps_1m

	ATLAS	ALICE	CMS	LHCb
Beta*	1.00 m	10.00 m	1.00 m	3.00 m
Crossing Angle (urad)	-120(V)	-80(V)	120(H)	-250(H)
Spectrometer Angle (urad)		no_value(V)		no_value(H)
Beam Separation (mm)	0(H)	.3(H)	-.5(V)	-.08(V)
Expected Collisions per turn	1318	39	1318	1296

	ATLAS	ALICE	CMS	LHCb
BPTX: deltaT of IP (B1-B2)	-0.03 ns	-0.07 ns	-0.06 ns	-0.11 ns
Luminous size (x,y) in um	19.6,20.7	-999.0,-999.0	19.1,13.8	44.2,45.1
Luminous size (z) in mm	55.0	-999.0	44.5	52.6
Lumi Centroid (x,y) in um	-49.1,1056.8	-999.0,-999.0	156.4,-674.2	465.5,-14.1
Lumi Centroid (z) in mm	-7.3	-999.0	7.9	6.6
Luminous Tilt in urads	-6.11,-60.37	-999.00,-999.00	105.43,205.42	-63.43,32.75

RESULTS: PHYSICS 1/2

- ATLAS b-jet triggers (dependent on beam spot)
  - High up time, fast bootstrap at beginning of fill
- Provided a plethora of data for beam studies
  - Follow position with sub micron statistical uncertainty

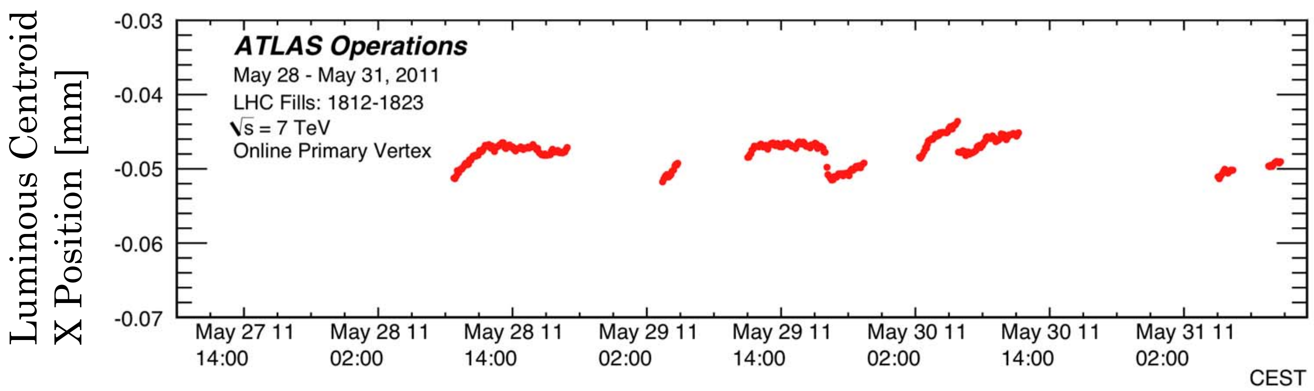

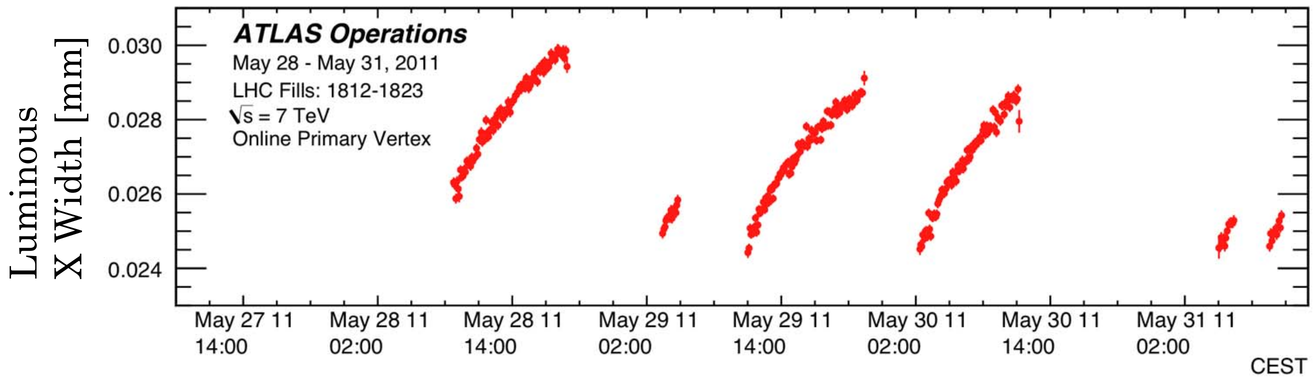

RESULTS: PHYSICS 2/2

- Measure position and width of each bunch (>1300)
- Needs high rate and devoted resolution calculation
- See unambiguous effects of beam-beam kicks on orbit

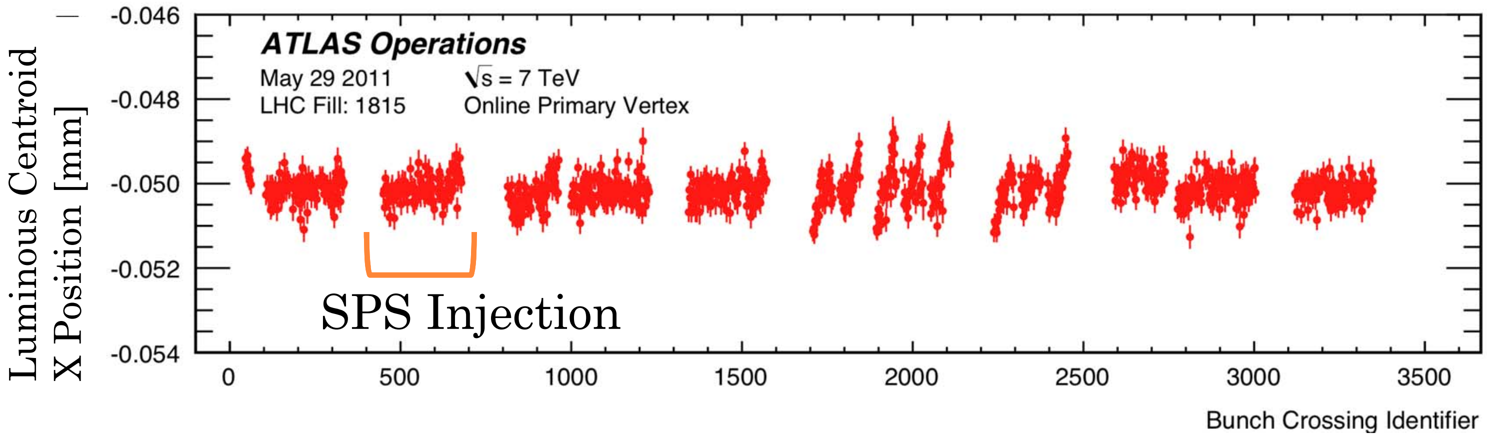

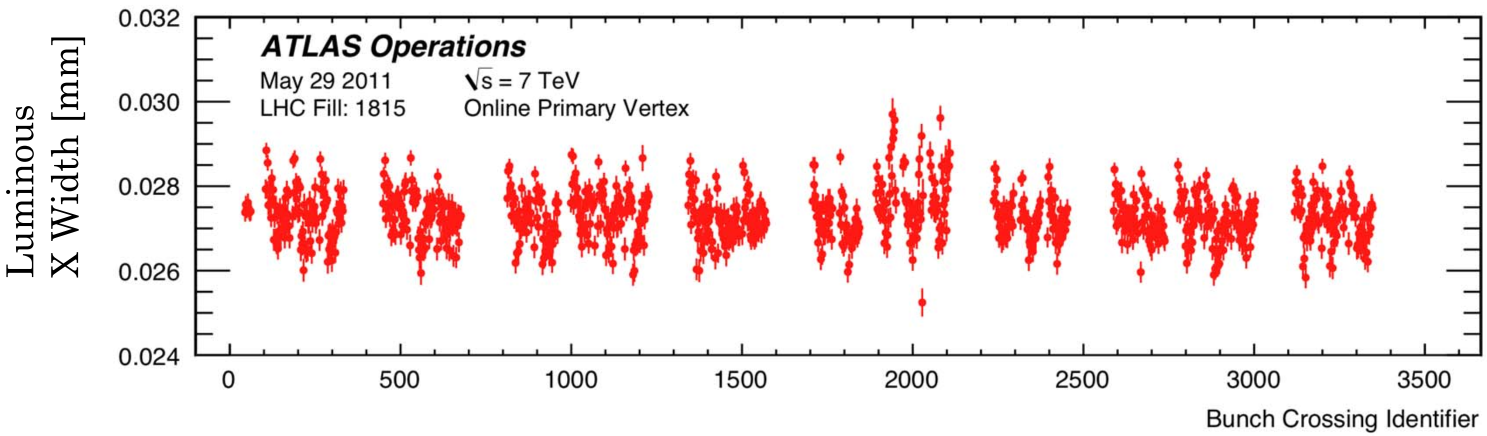

CONCLUSION

- Built a system to measure the beam spot
  - On the HLT in **near real** time with large rate
  - Measure position with < $1\,\mu m$ statistical uncertainty
- Feedback the answer to 13,000 processes
  - Sharp change across the LumiBlock boundary
  - So fast to update with **proxies, with no DAQ busy!**
  - Tracks beam parameter drifts within $2\,\mu m$
- Provide new data for LHC development
  - Extremely **accurate per bunch** measurements
  - Trending during runs, after long stops, etc …

BACK UPS

- Many distributions of LHC beam parameters

VERTEX DISTRIBUTIONS

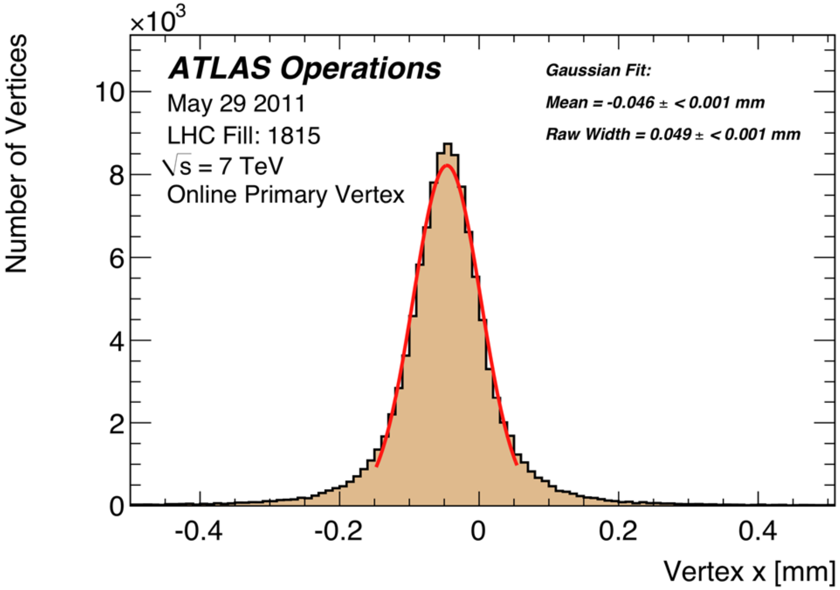

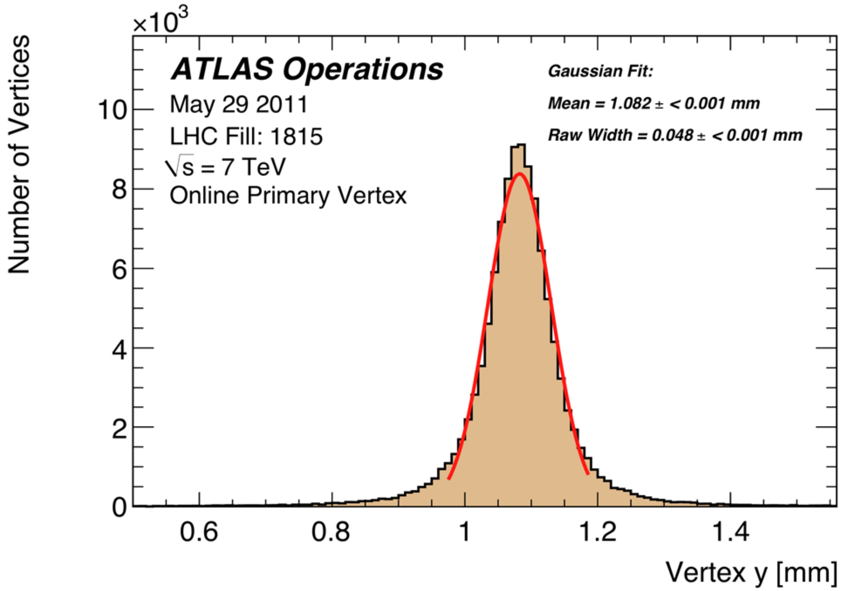

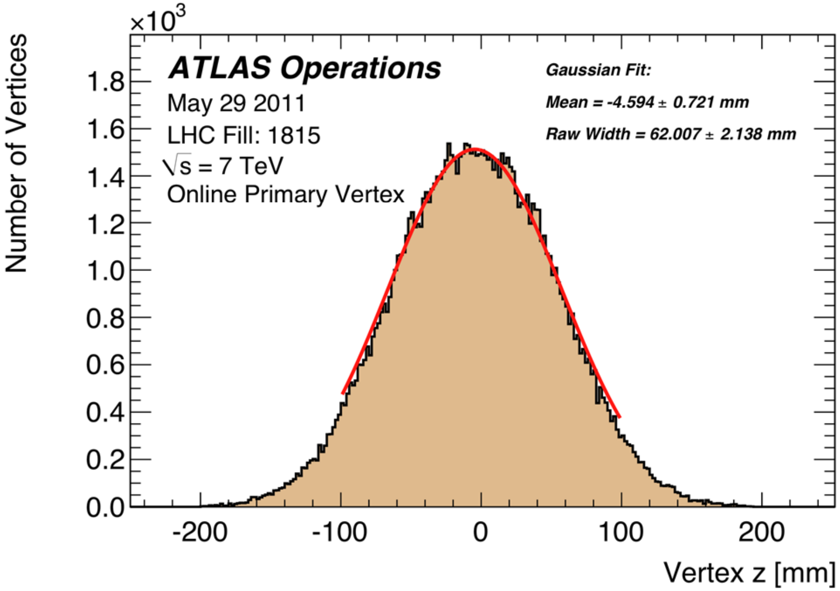

TILTS

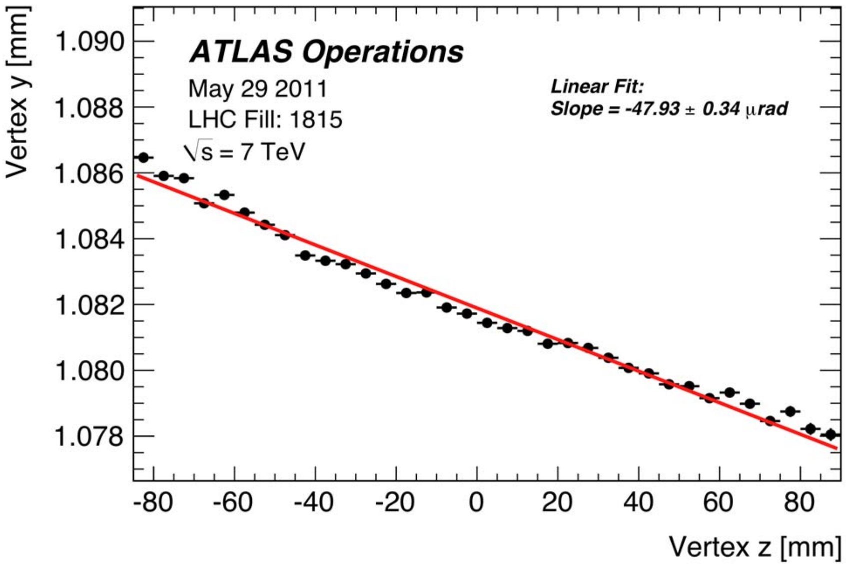

ATLAS Operations
May 29 2011
LHC Fill: 1815
√s = 7 TeV
Linear Fit:
Slope = -47.93 ± 0.34 μrad
Vertex y [mm]
Vertex z [mm]

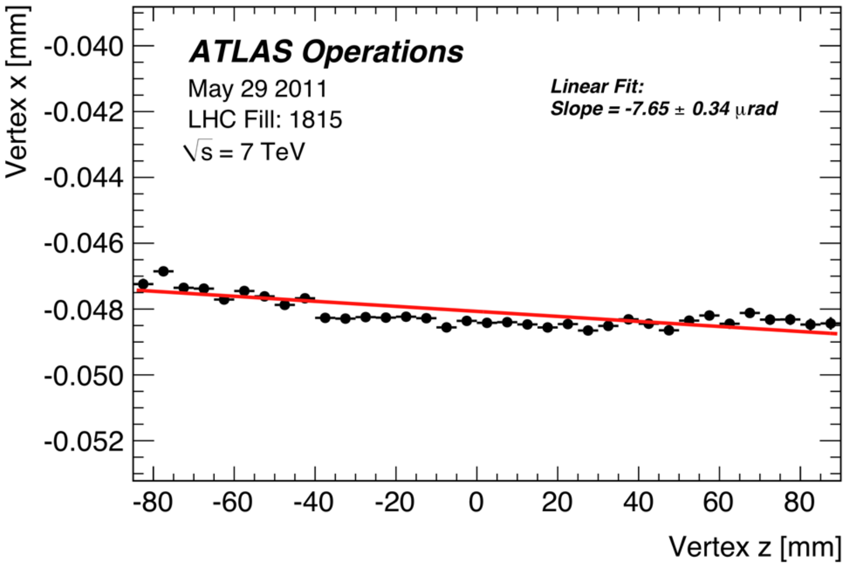

ATLAS Operations
May 29 2011
LHC Fill: 1815
√s = 7 TeV
Linear Fit:
Slope = -7.65 ± 0.34 μrad
Vertex x [mm]
Vertex z [mm]

D0 VS Φ

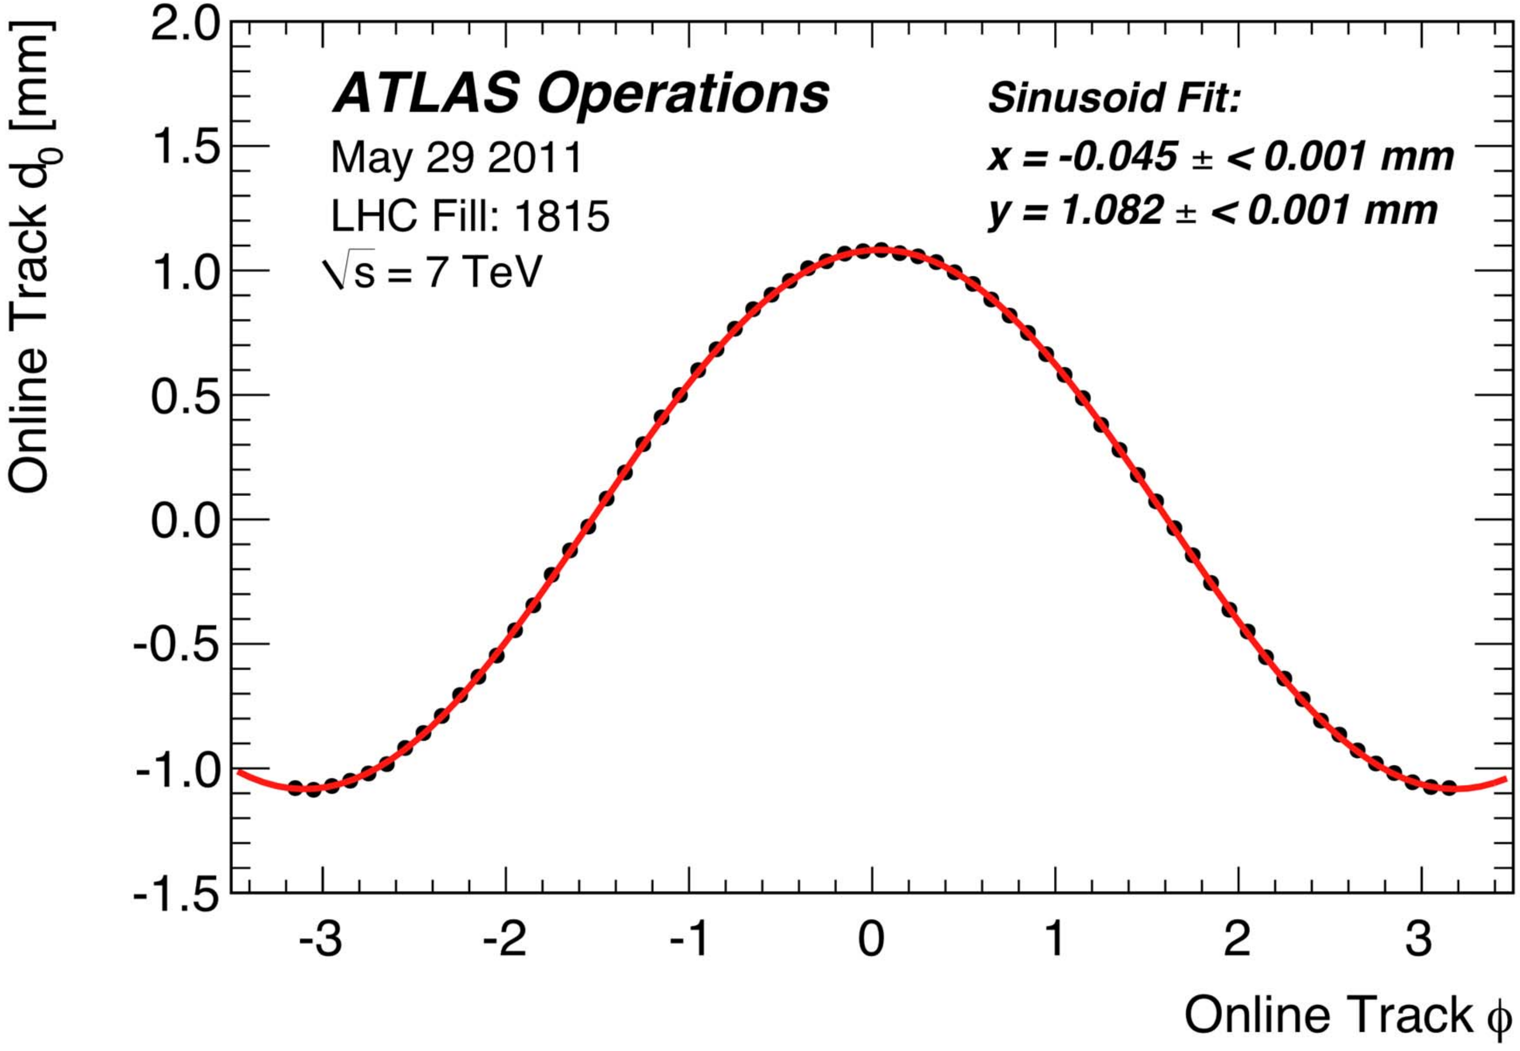

ATLAS Operations
May 29 2011
LHC Fill: 1815
√s = 7 TeV
Sinusoid Fit:
x = -0.045 ± < 0.001 mm
y = 1.082 ± < 0.001 mm
Online Track d0 [mm]
Online Track ϕ
2.0
1.5
1.0
0.5
0.0
-0.5
-1.0
-1.5
-3
-2
-1
0
1
2
3

SPLIT VERTEX RESOLUTION VS NUMBER OF TRACKS

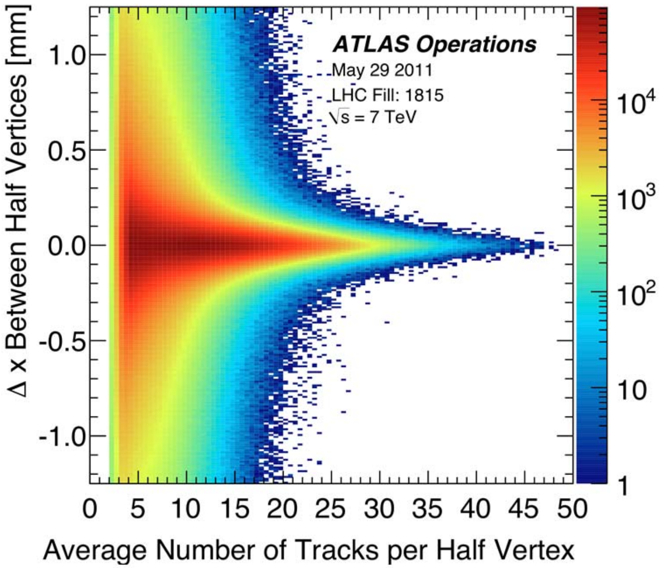

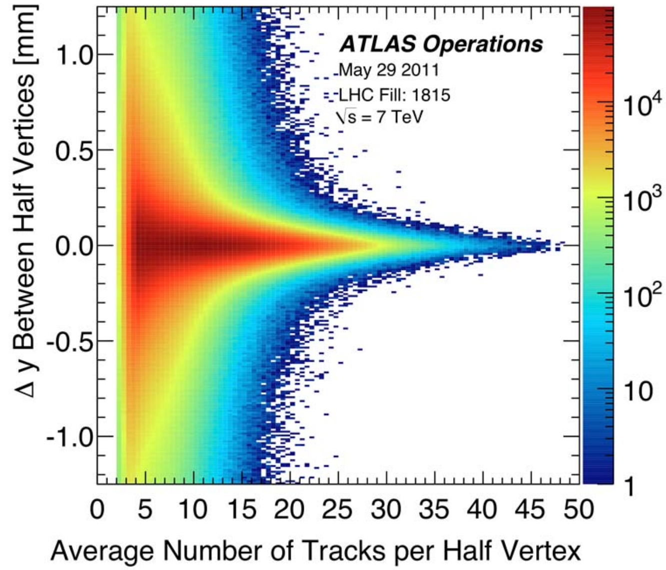

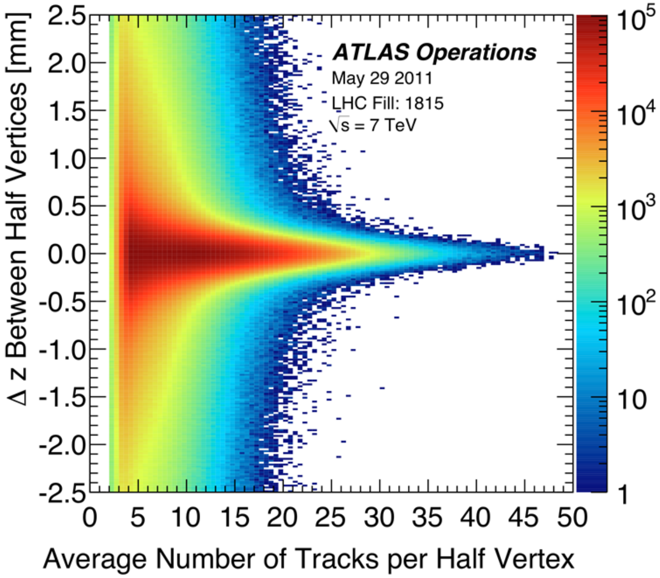

RESOLUTION VS NUMBER OF TRACKS

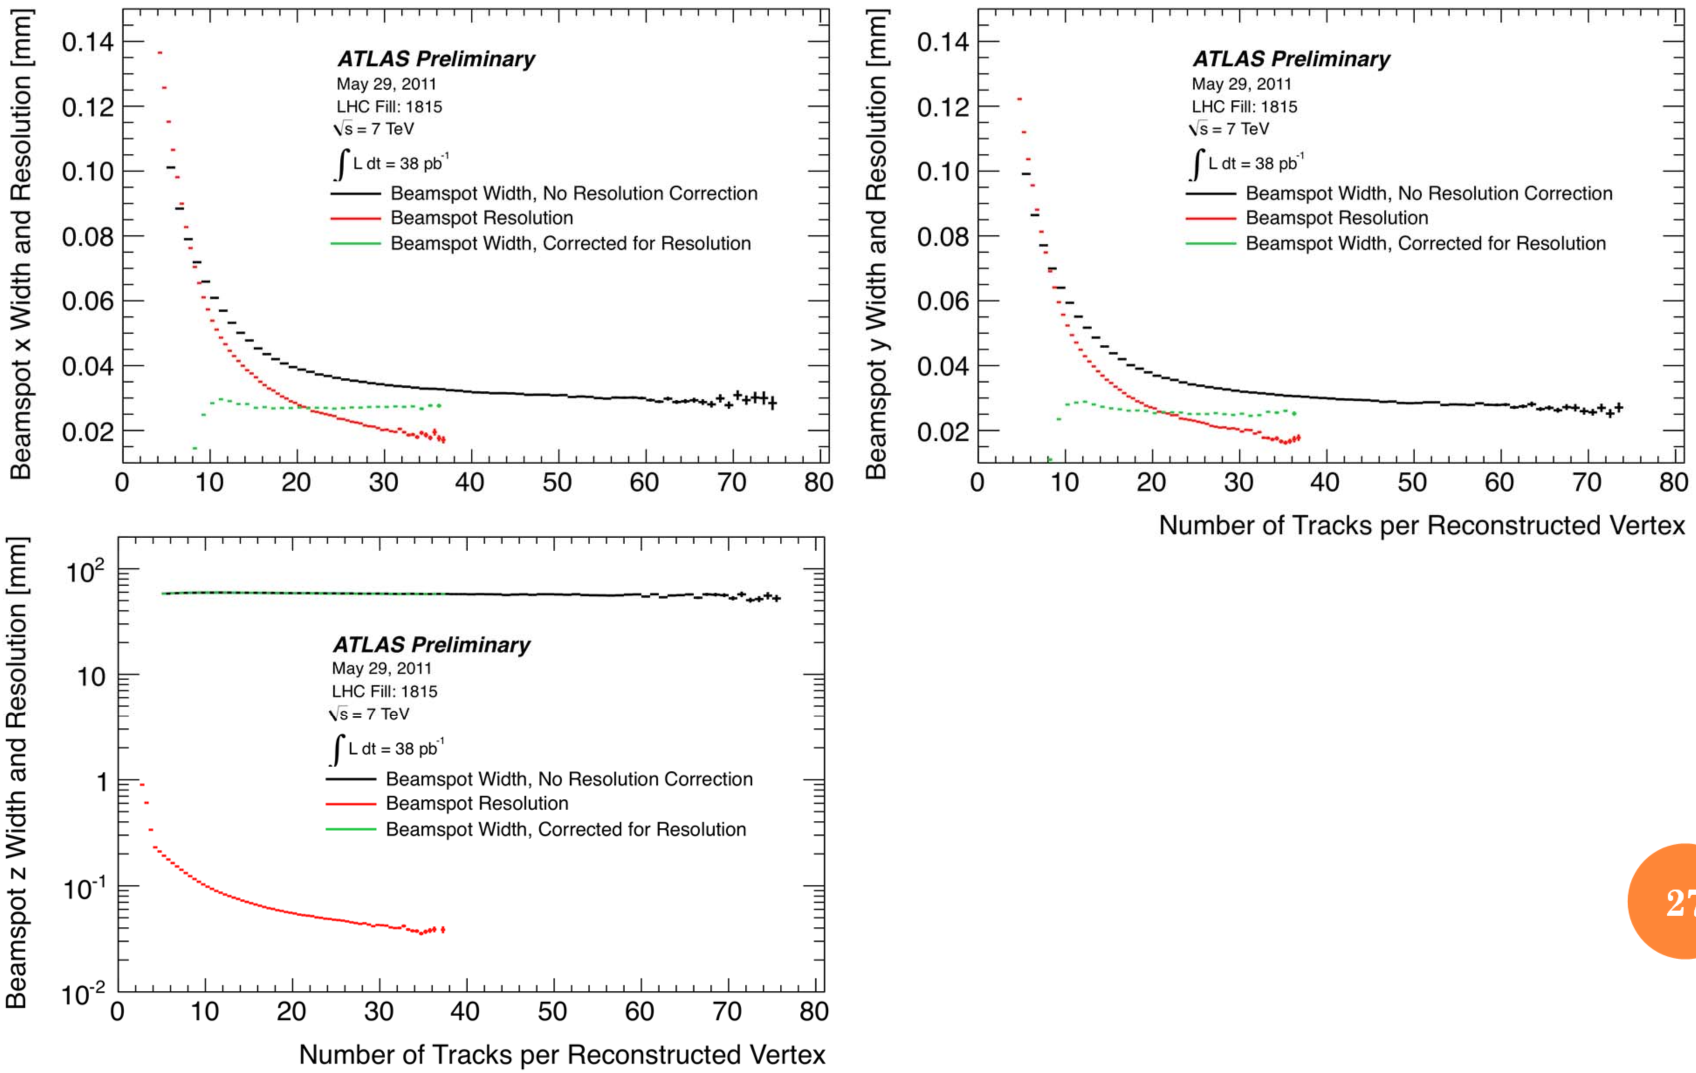

POSITION VS TIME

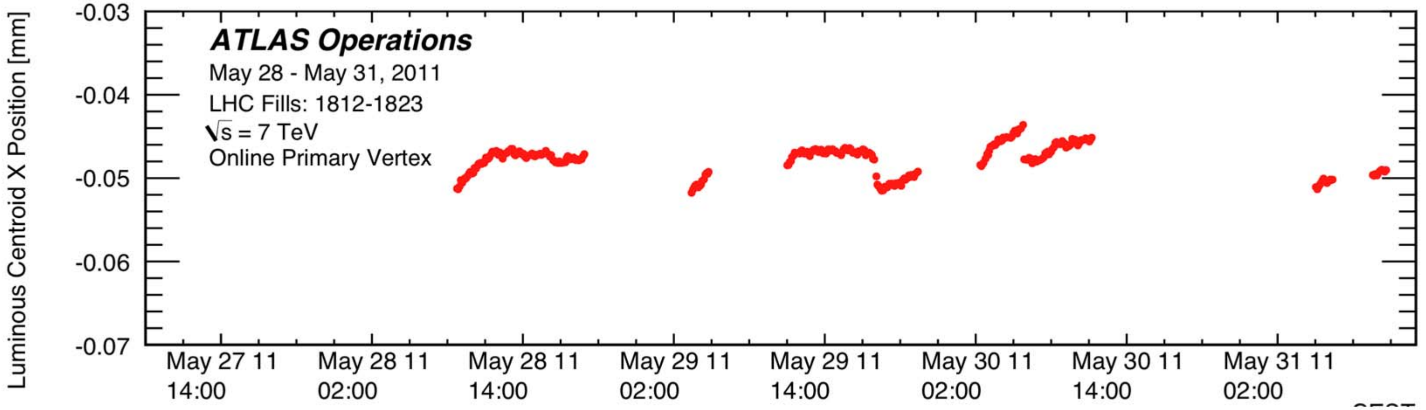

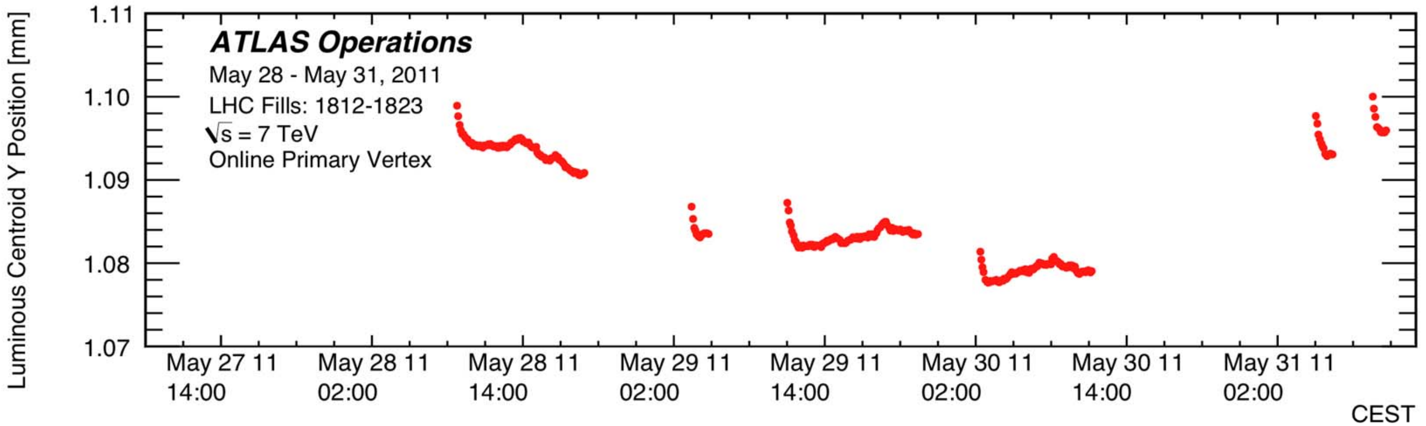

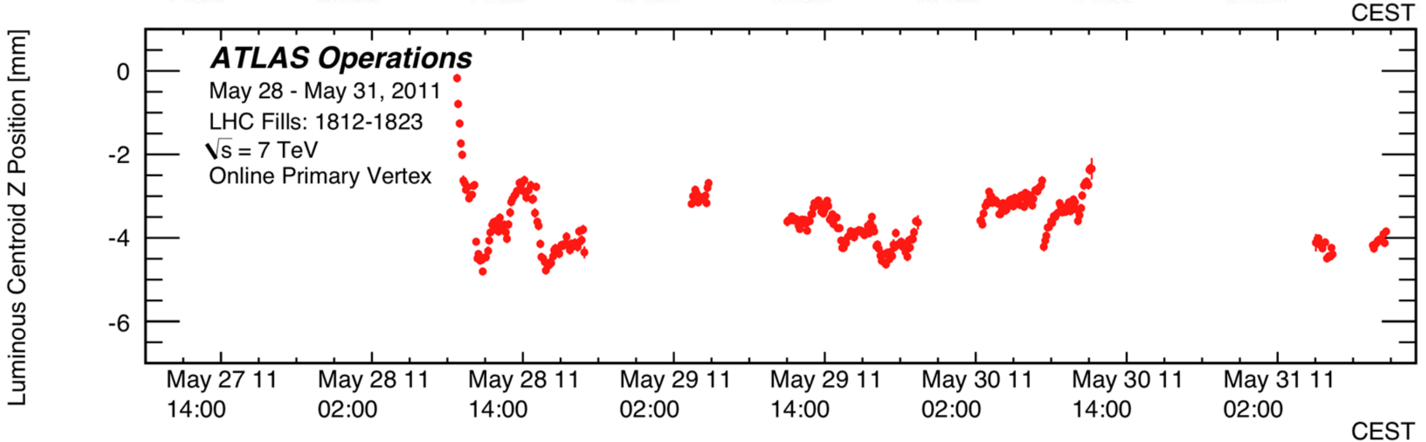

WIDTH VS TIME

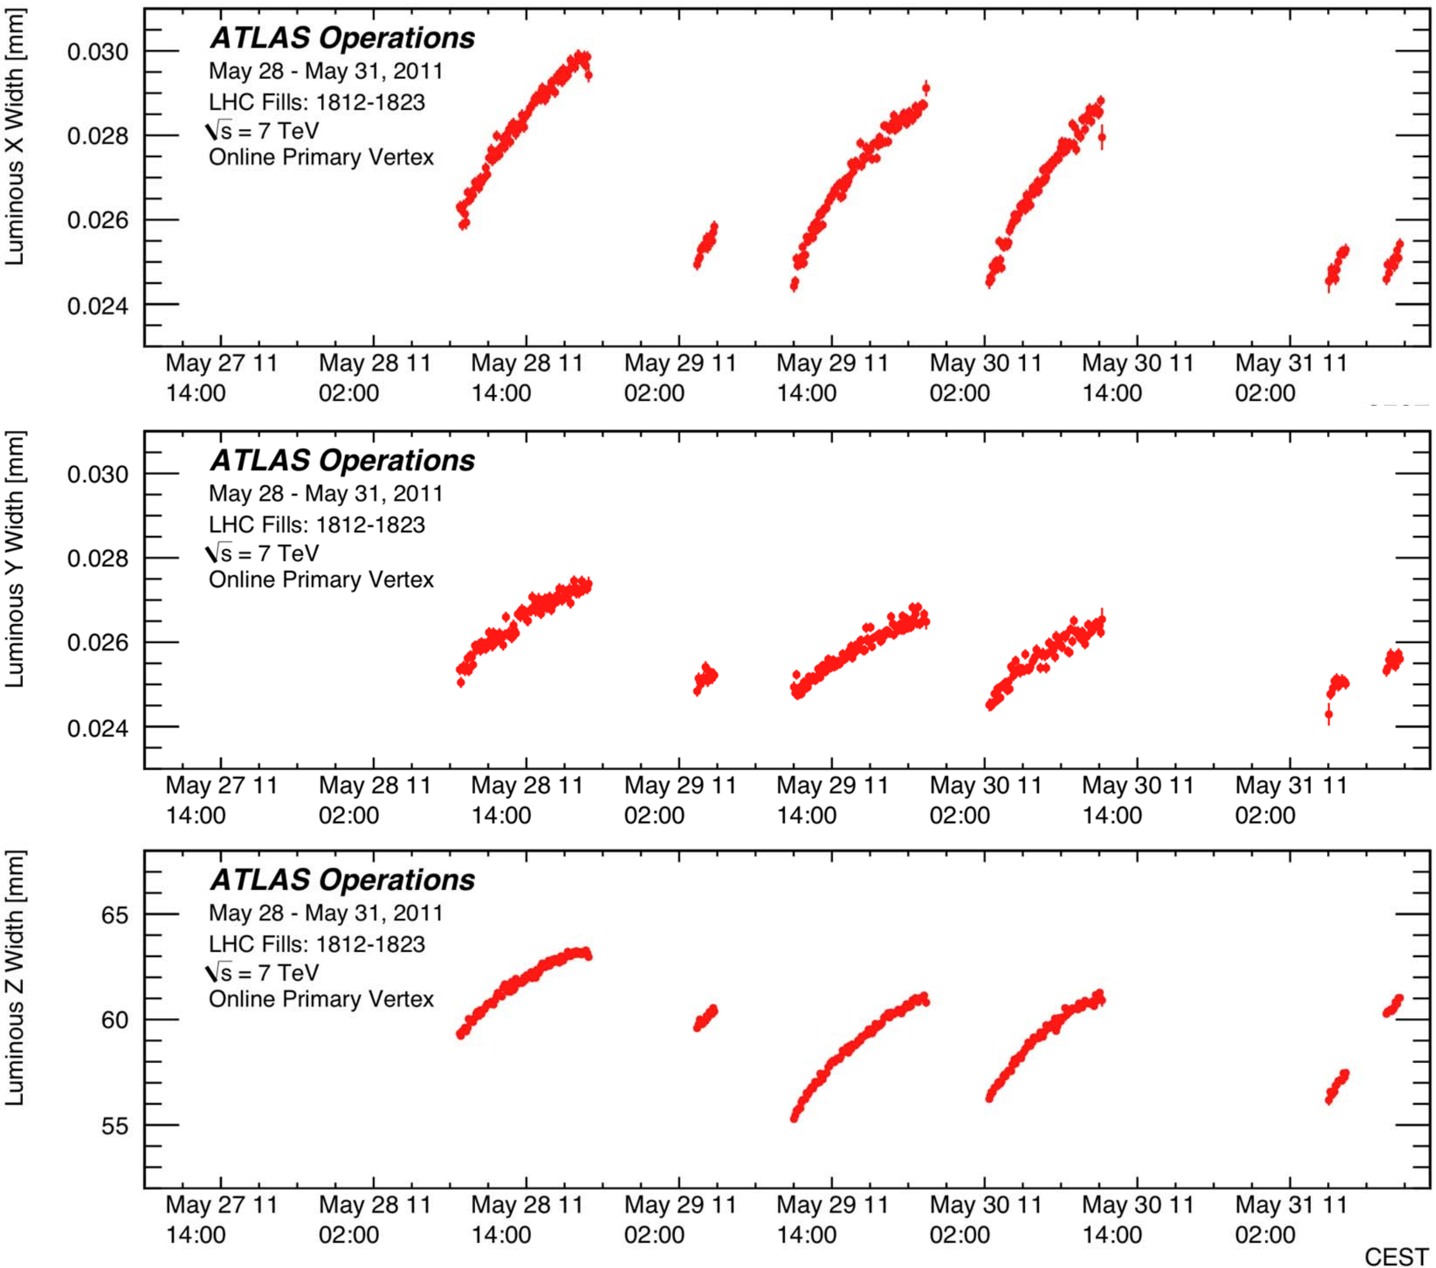

PER BUNCH POSITIONS

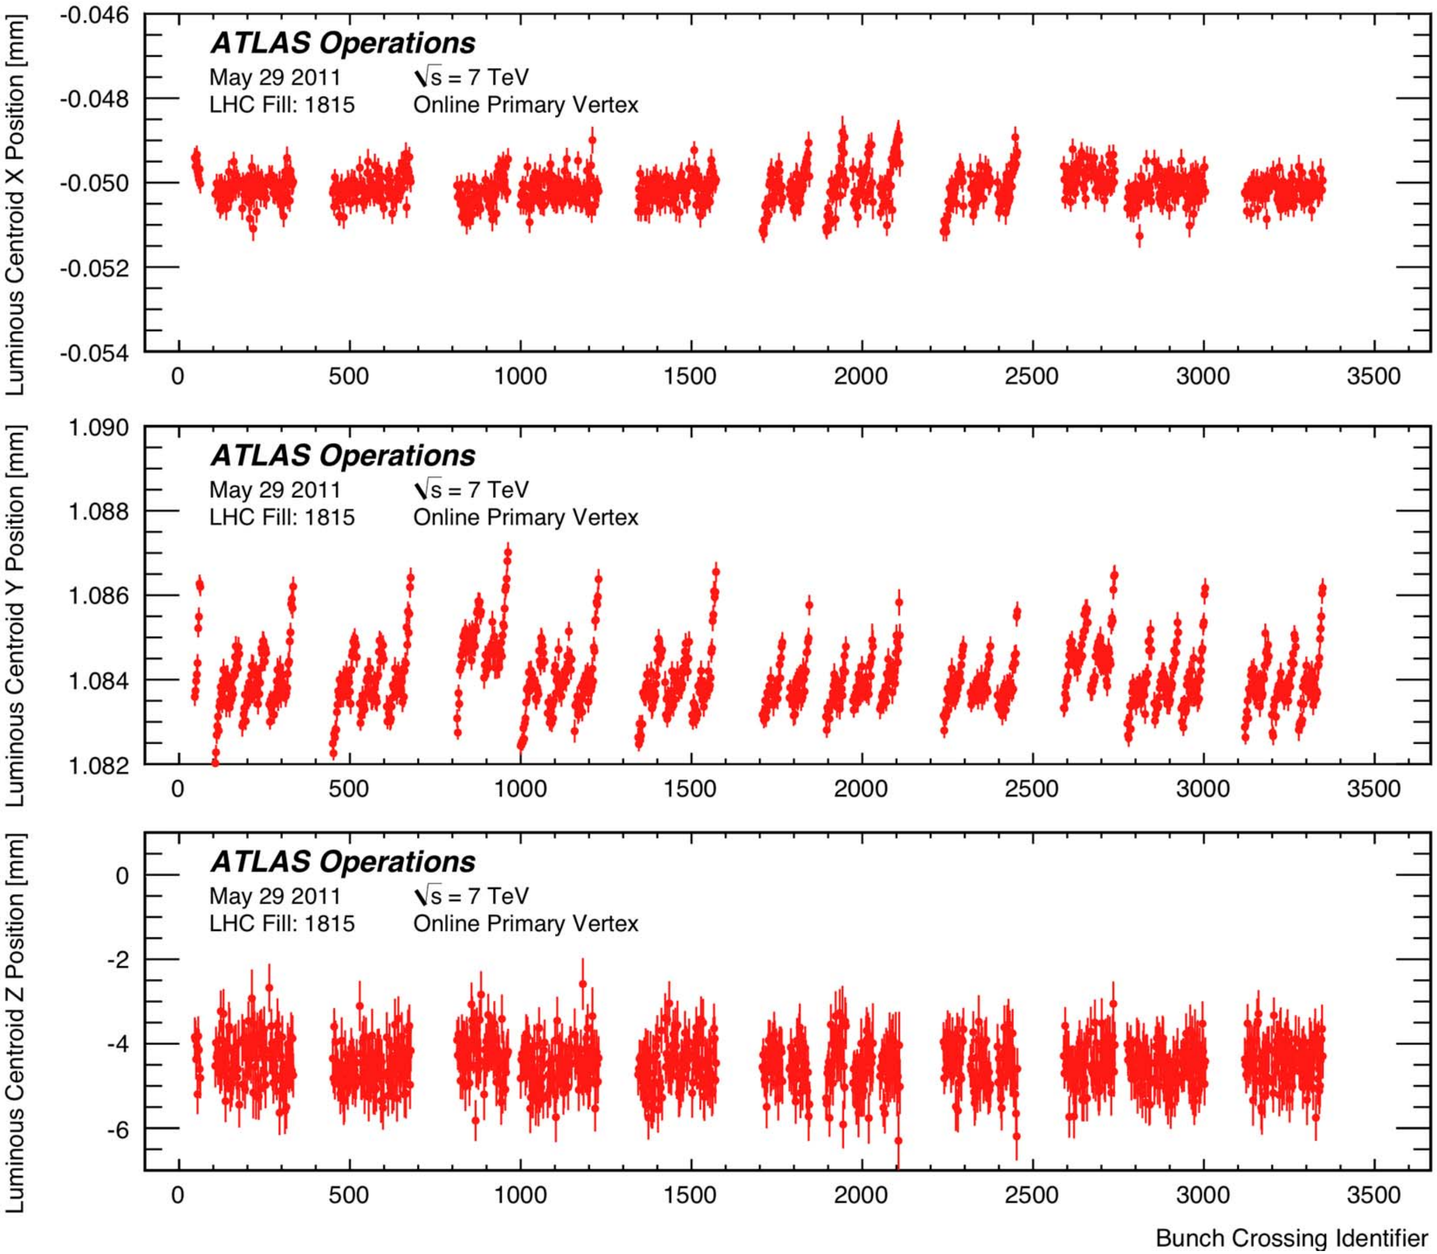

PER BUNCH WIDTHS

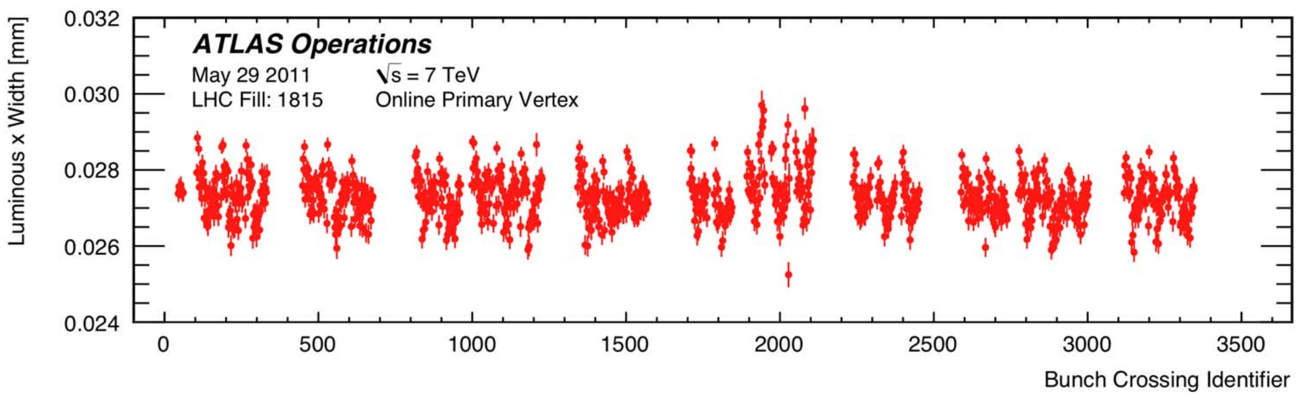

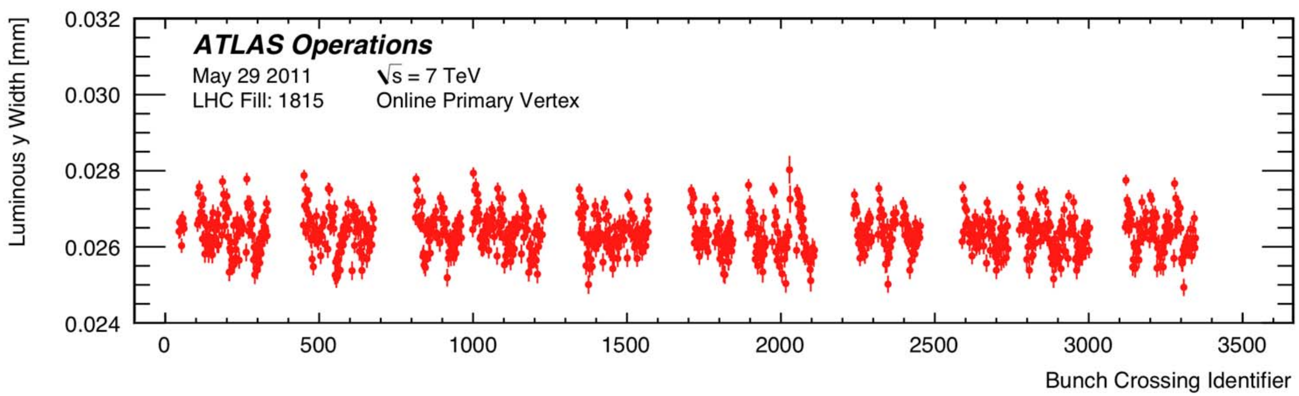

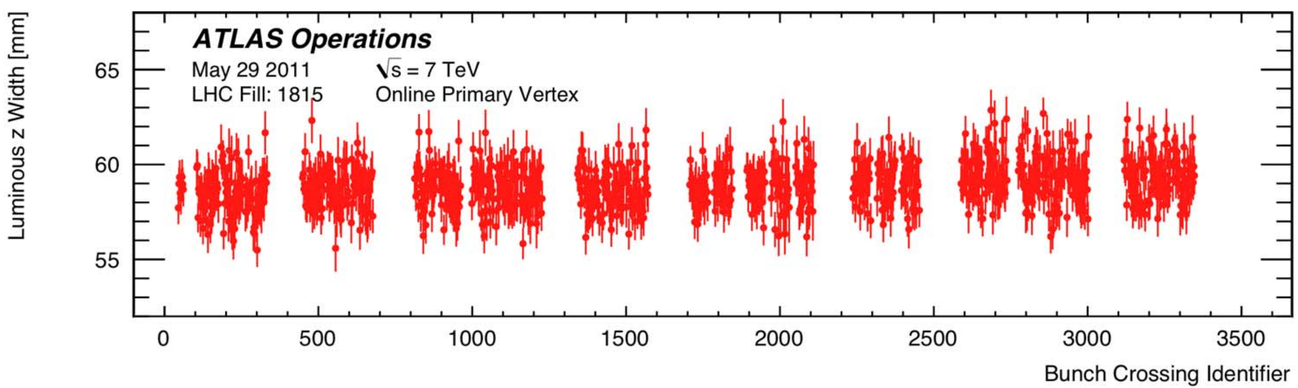

WHAT'S THE PROBLEM?

- ATLAS produces crap-tons of data!
- Wait, why?
  - Collisions are usually pretty "boring"

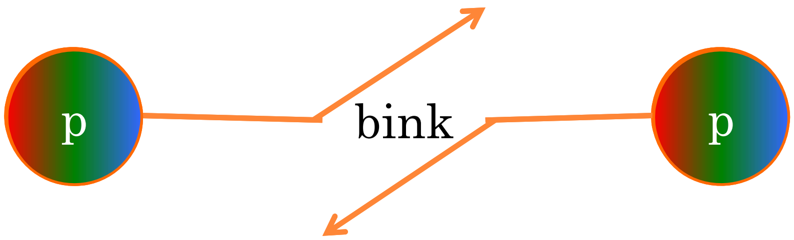

- Low energy processes already studied
  - Inelastic scattering
  - Dijet production
  - W/Z/γ
  - …
- Yesterday's signal is today's background and tomorrow's noise

WHY SO MUCH DATA?

- Only rarely does something “interesting” happen

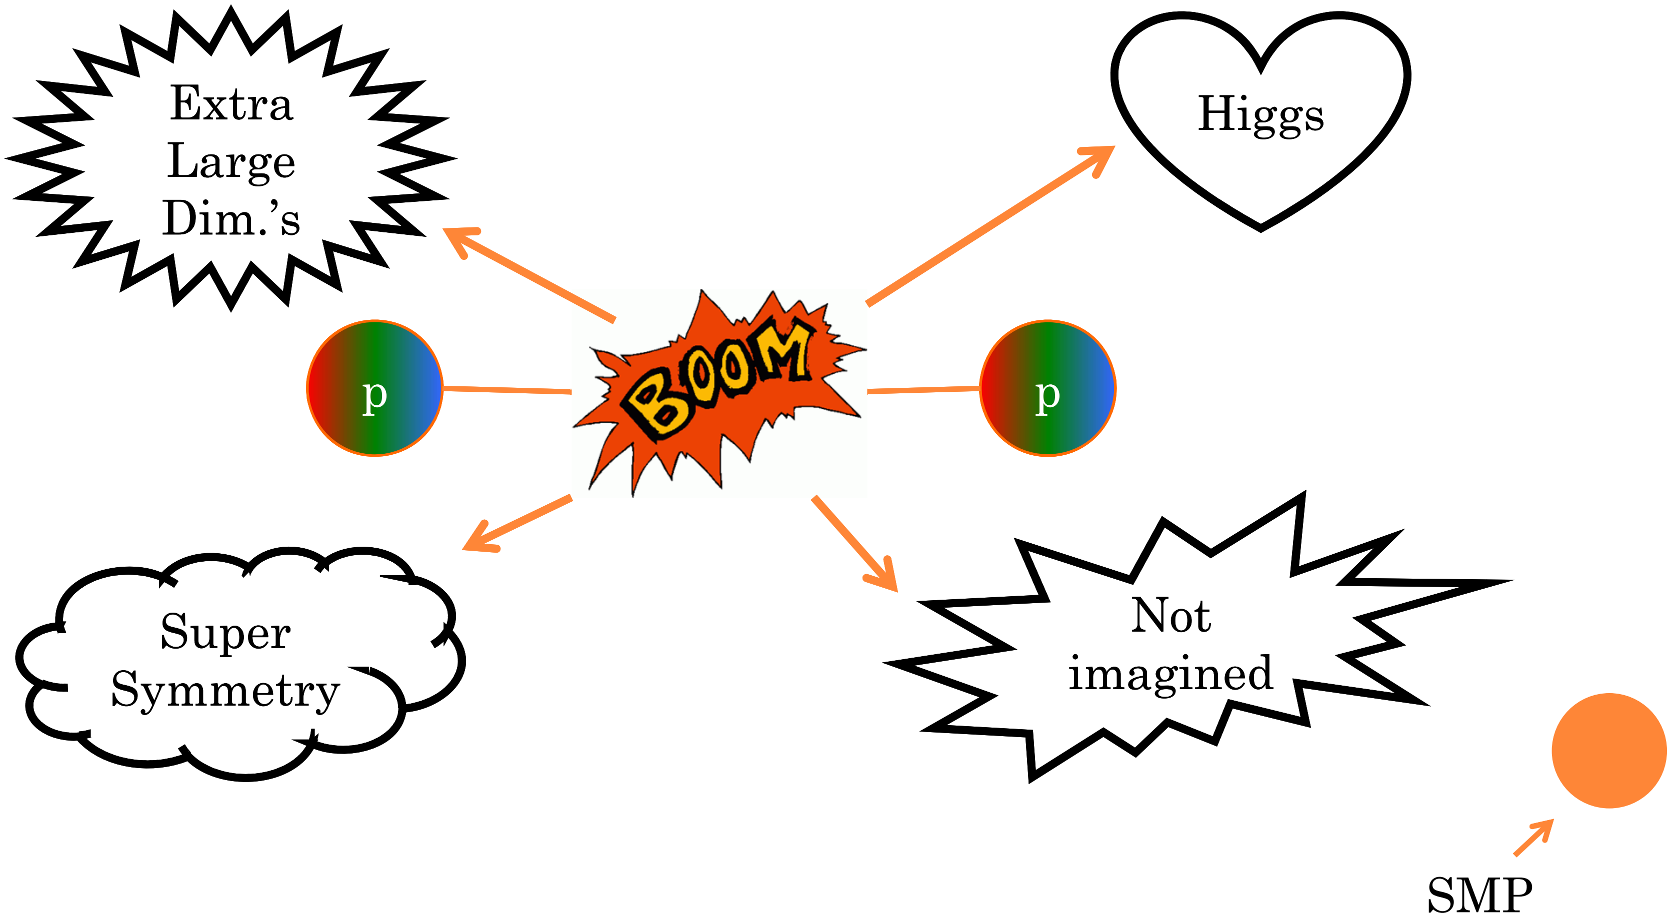

CONTEXT

- ATLAS one of several large detectors at LHC
- LHC delivers ~15 Million bunch crossing/second
  - Most collisions are "boring" and can be thrown out
  - Rare few could be a Higgs, black hole, SUSY etc.
- Recording all the data would be **20 TB/second!**
  - Need to **trigger** data acquisition on interesting events