The PHENIX Experiment in the RHIC Run 7

(with some references to Run 6...)

Martin L. Purschke, Brookhaven National Laboratory
for the PHENIX Collaboration

RHIC from space

Long Island, NY

Or, much simpler:

Our best run ever!

RHIC/PHENIX at a glance

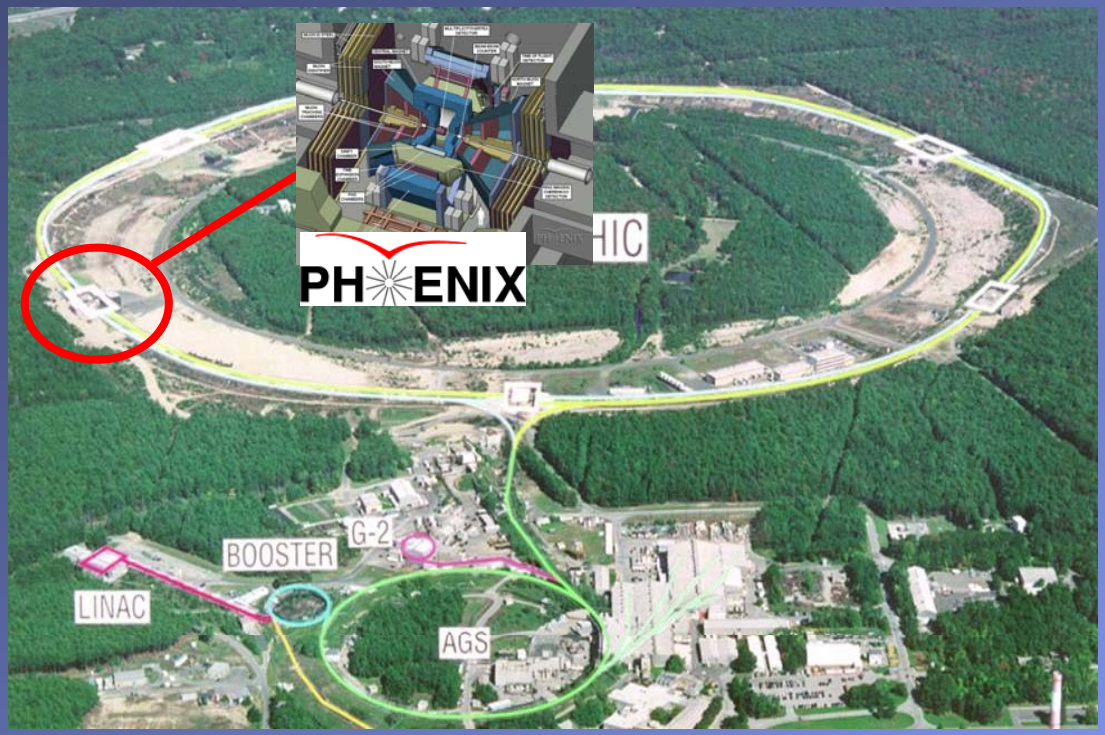

RHIC:

2 independent rings, one beam clockwise, the other counterclockwise

sqrt(S_{NN})= 500GeV * Z/A

~200 GeV for Heavy Ions

~500 GeV for proton-proton (polarized)

PHENIX:

4 spectrometer arms

15 Detector subsystems

500,000 detector channels

Lots of readout electronics

Uncompressed Event size typically 280 -220 - 110 KB for AuAu, CuCu, pp

Data rate ~5KHz (Au+Au)

Front-end data rate 0.5 - 1.1 GB/s

Data Logging rate ~400MB/s, 700 MB/s max

...and 4 new Detector Systems in Run 7

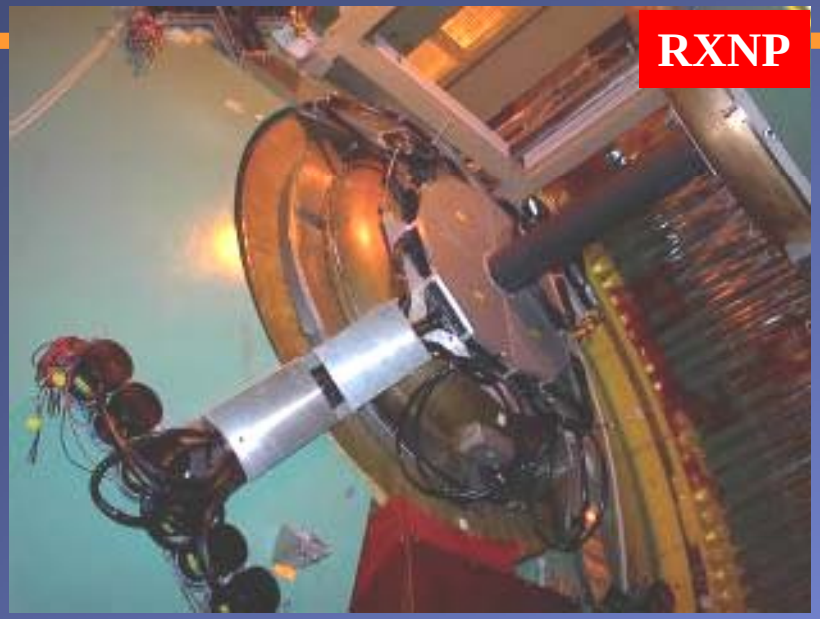

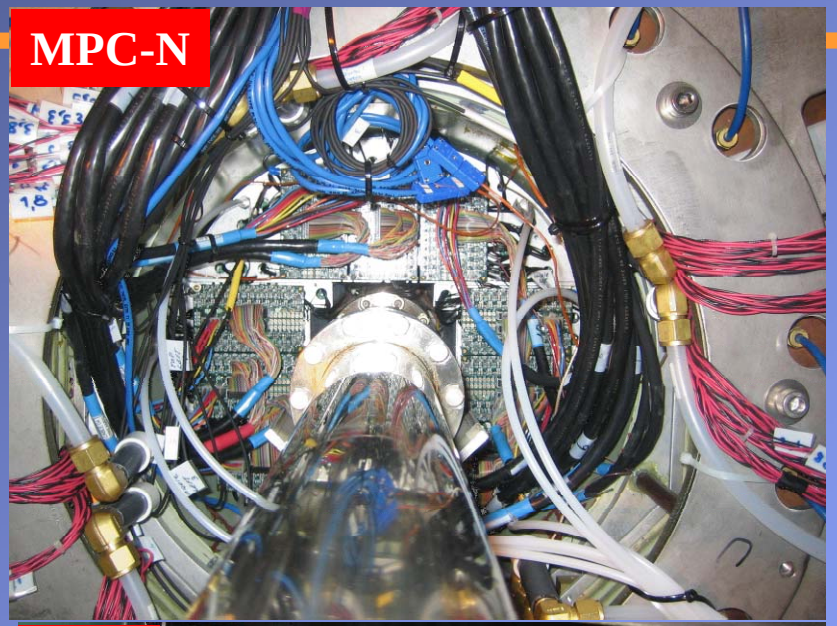

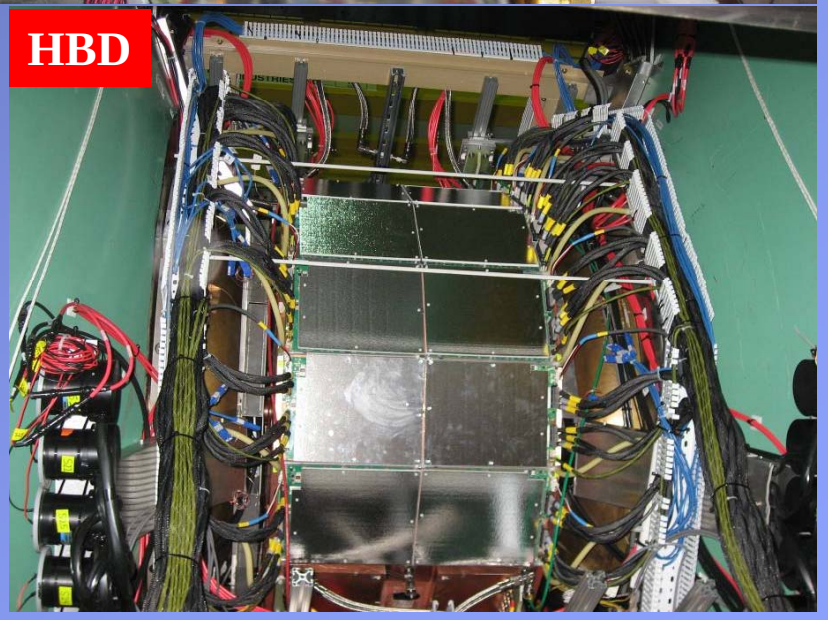

Building up to record speed

- Over the previous runs we have been adding improvements
- Had lighter systems, d+Au, p-p, Cu-Cu in the last runs, less of a challenge than 200GeV Au+Au

Ingredients:

- Distributed data compression (run 4)
- Multi-Event buffering (run 5)
- Mostly consolidating the achievements/tuning/etc in run 6, also lots of improvements in operations (increased uptime)
- 10G Network upgrade in run 7, added Lvl2 filtering

With increased luminosity, we saw the previously demonstrated 600++MB/s data rate in earnest for the first time.

Data Compression

Found that the raw data are still gzip-compressible after zero-suppression and other data ***reduction*** techniques

Introduced a compressed raw data format that supports a late-stage compression

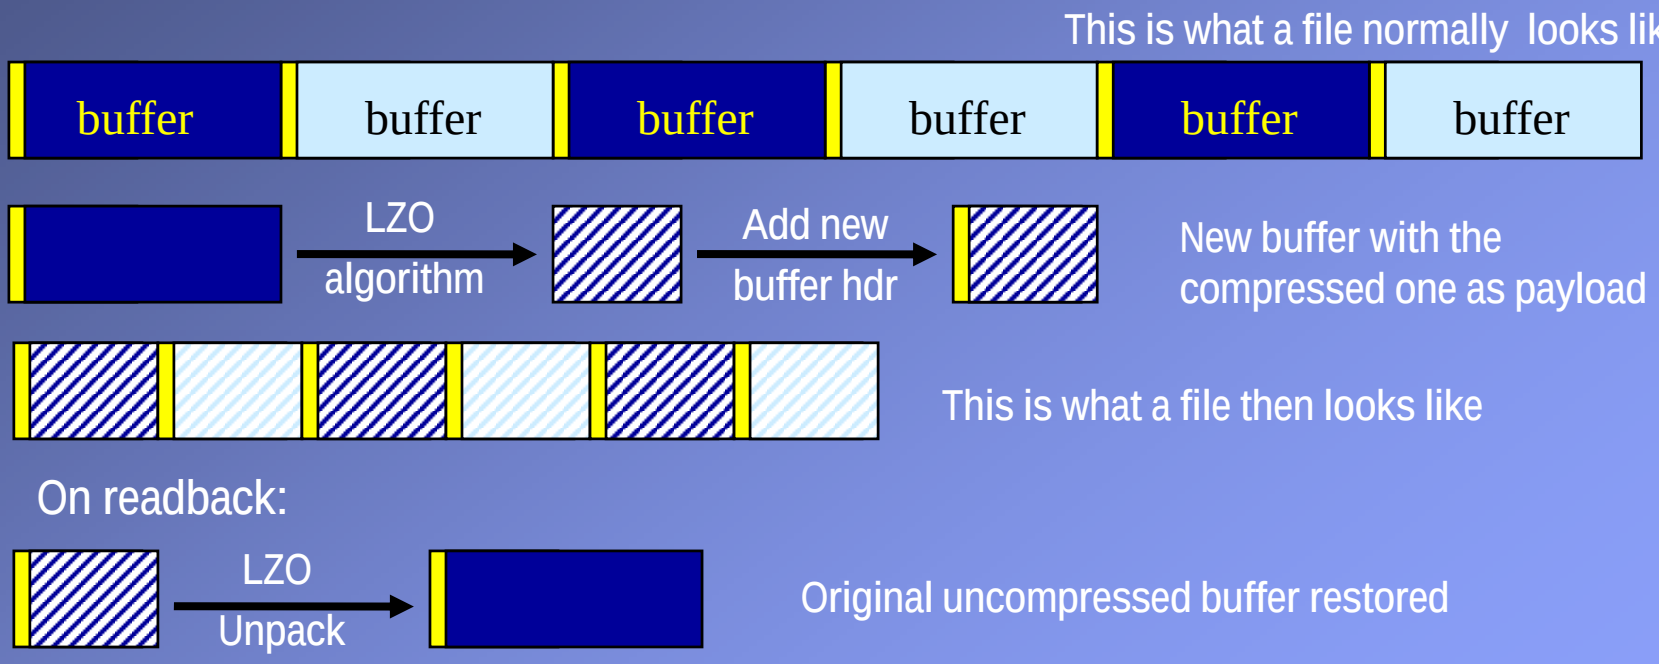

All this is handled completely in the I/O layer, the higher-level routines just receive a buffer as before.

Distributed Compression

The Event builder has to cope with the uncompressed data flow, e.g. 600MB/s … 1200MB/s

The compression is handled in the "Assembly and Trigger Processors" (ATP's) and can so be distributed over many CPU's -- that was the breakthrough

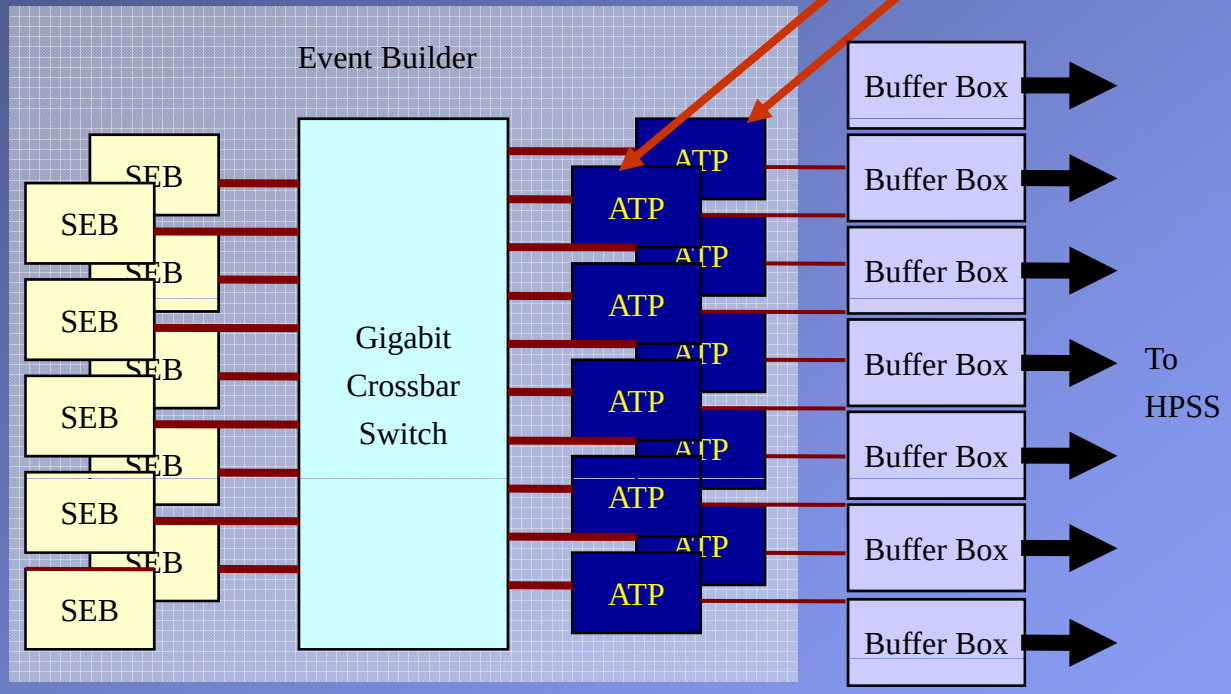

The buffer boxes and storage system see the compressed data stream, 350MB/s … 650MB/s

Multi-Event Buffering: DAQ Evolution

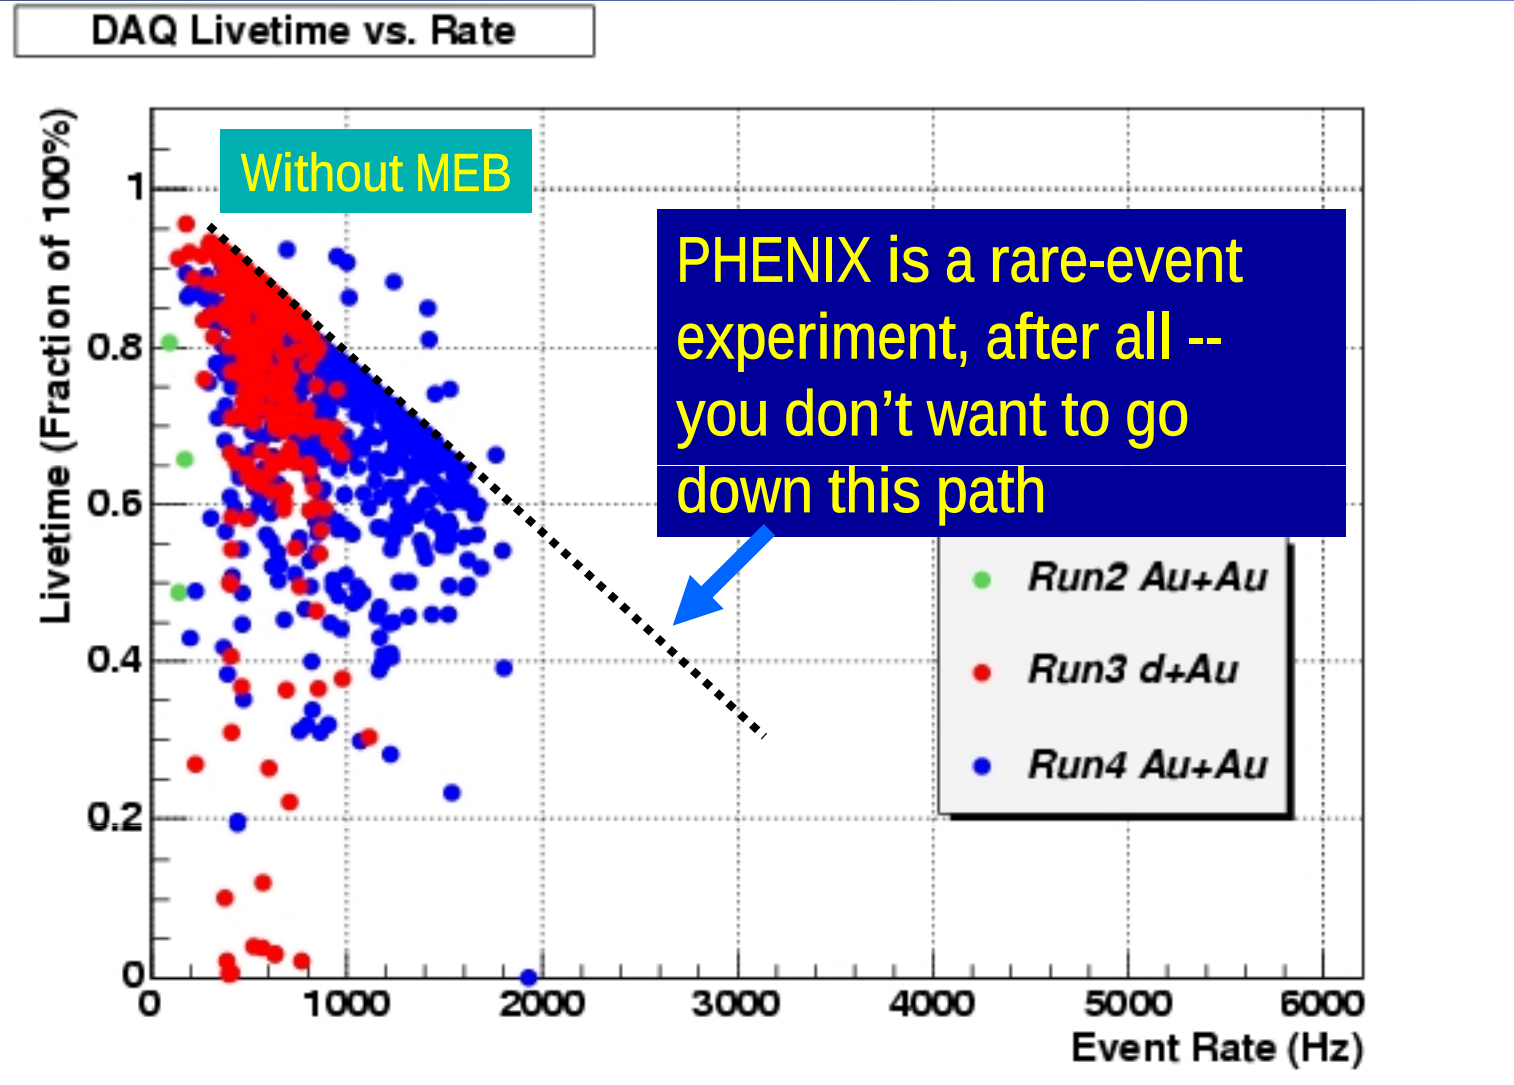

MEB: trigger delays by analog Memory

trigger electronics needs to buy some time to make its decision

done by storing the signal charge in an analog memory (AMU)

Memory keeps the state of some 40us worth of bunch crossings

Trigger decision arrives. FEM goes back a given number of analog memory cells and digitizes the contents of that memory location

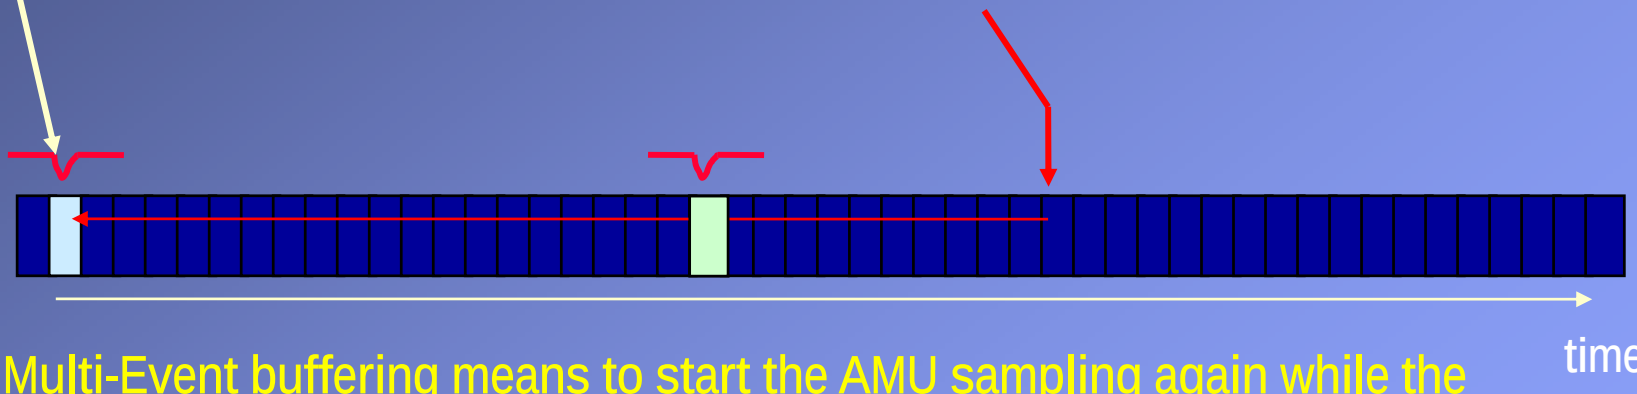

Multi-Event buffering means to start the AMU sampling again while the current sample is still being digitized.

Trigger busy released ***much*** earlier

deadtime greatly reduced

The Multi-Event Buffering Effect

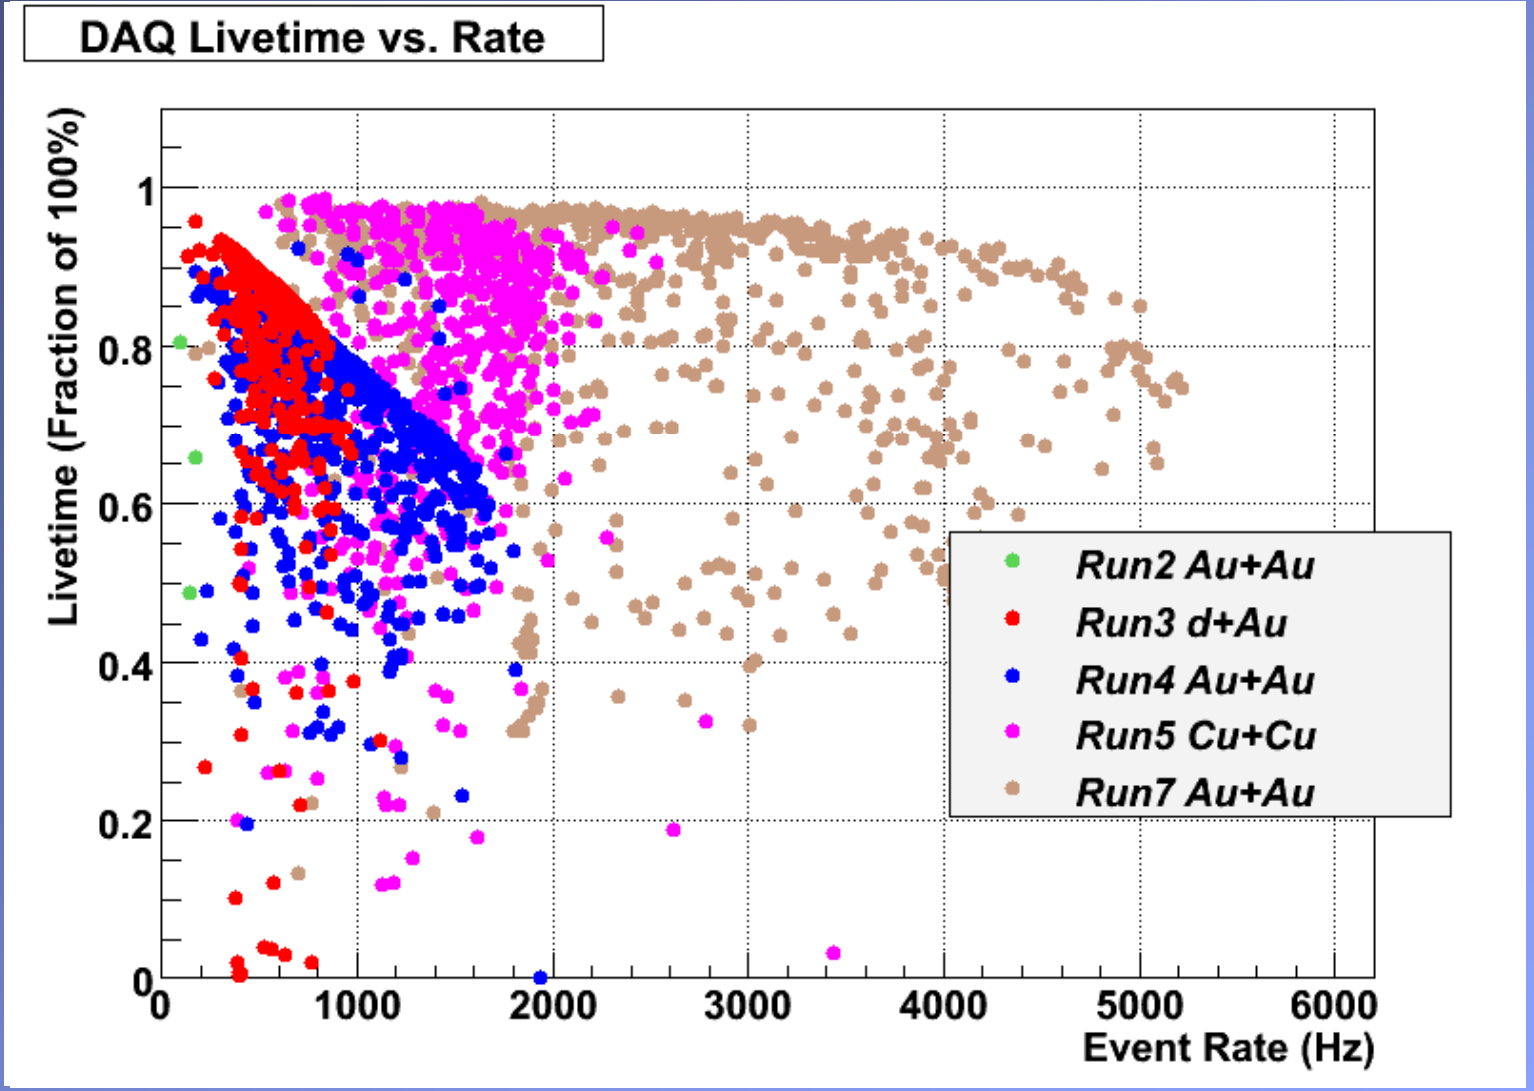

~600 MB/s

This shows the aggregated data rate from the DAQ to disk in a RHIC fill
We are very proud of this performance...

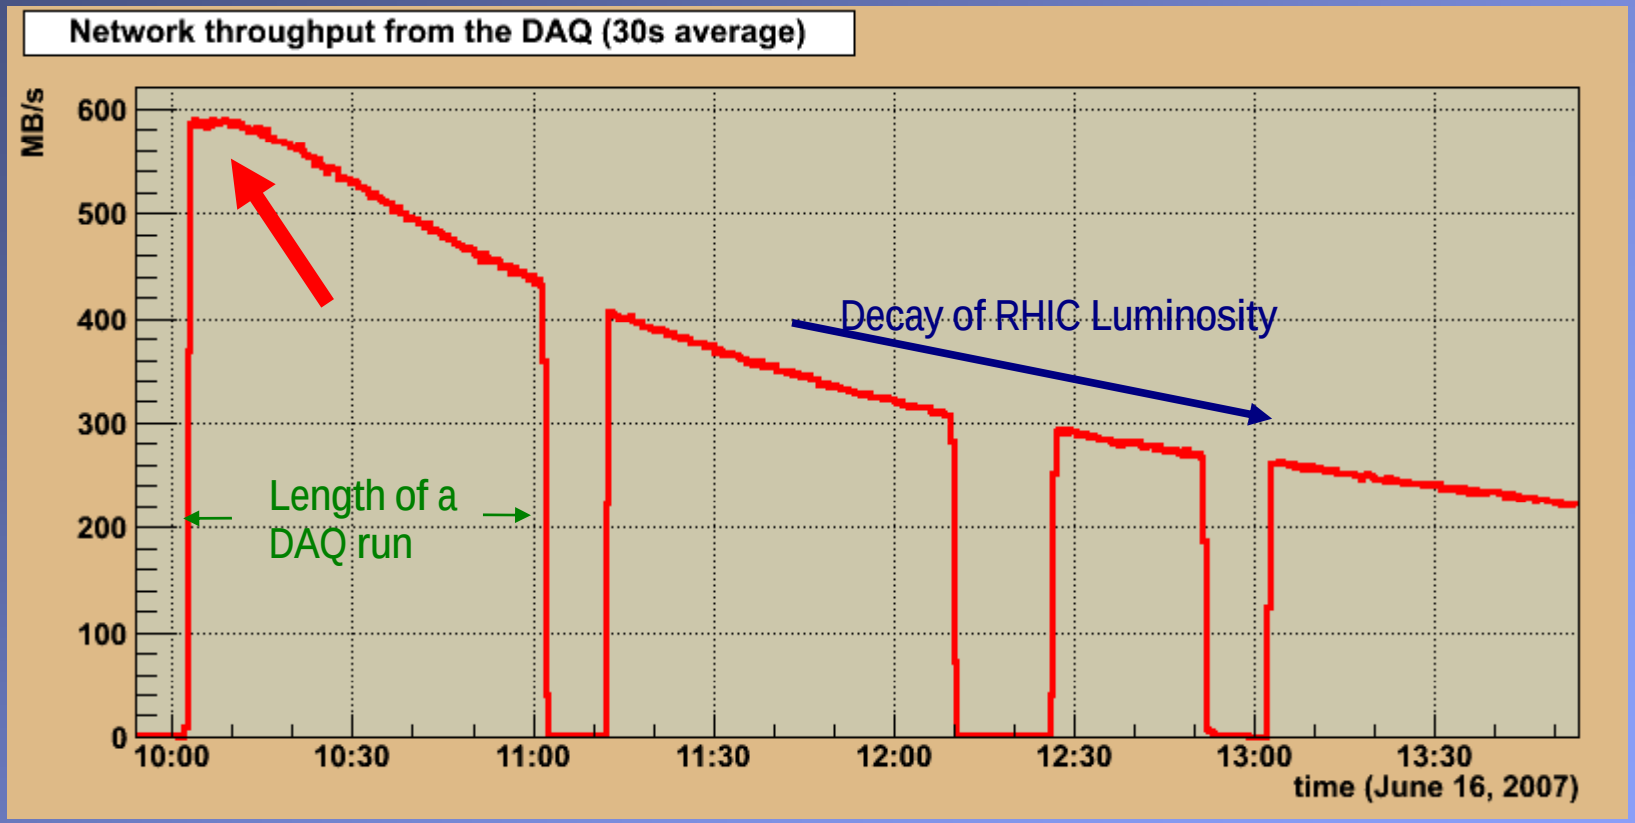

It's not the best, it's one where I was there...
best went up to 650MB/s

Event statistics

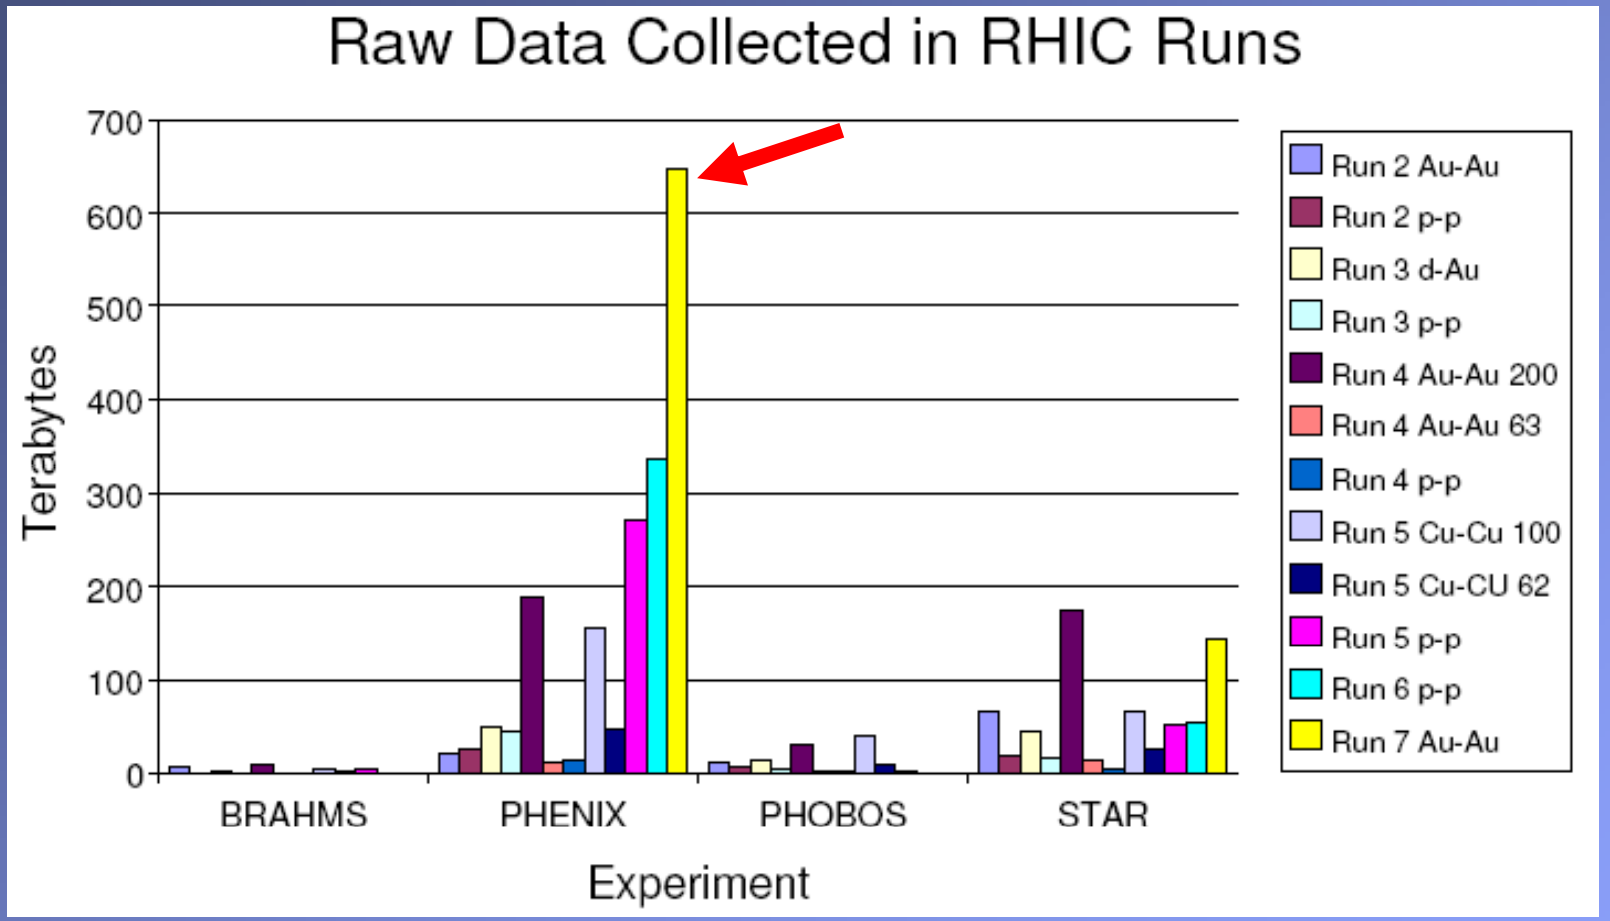

5.7 Billion Events in ~650TB of data

Run 6 pp – 6.8 Billion @200GeV
1 Billion @62GeV

Online Filtering and Reconstruction

- We ran Level-2 triggers in the ATP's in so-called filter mode
- lvl2 triggers don't reject but fish out interesting events for priority reconstruction
- Filtered data were sent to IN2P3 in France where resources were available AND where the people most interested in the filtered dataset are
- Invaluable tool to get a reading how you are doing, as well as preliminary physics signals
- Used to refine our GRID file transfer procedures to "new" remote sites (not that much data volume transferred during this run, ~50TB – Run6 - 300TB)

Summary

- Very successful run, 650TB of data on tape
- Can do > 600MB/s
- 4 new detector systems which still needed some “shakedown”
- Reached 5KHz event rate in Au-Au despite larger event size
- successful filtering effort for priority reconstruction

Where we are w.r.t. others

All in MB/s
all approximate

~25

LHCb

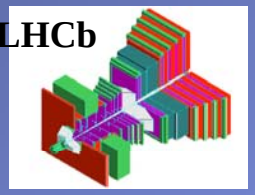

~40

~100

CMS

~100

~150

ATLAS

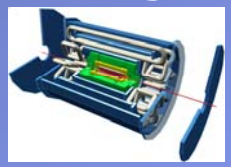

~300

PHENIX

600

ALICE

~1250

400-600MB/s are not
so Sci-Fi these days