The Readout and Data Acquisition Design of the sPHENIX Detector at RHIC

Martin L. Purschke

BROOKHAVEN NATIONAL LABORATORY

Manhattan

Long Island, NY

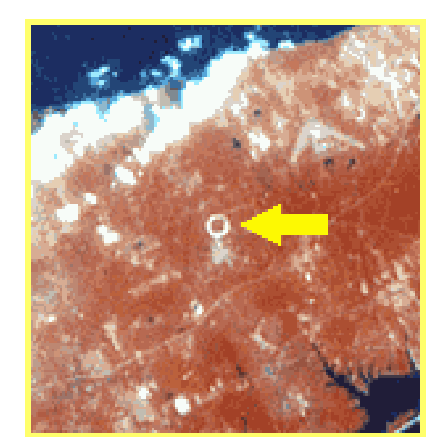

RHIC from space

PHENIX 2000 - 2016

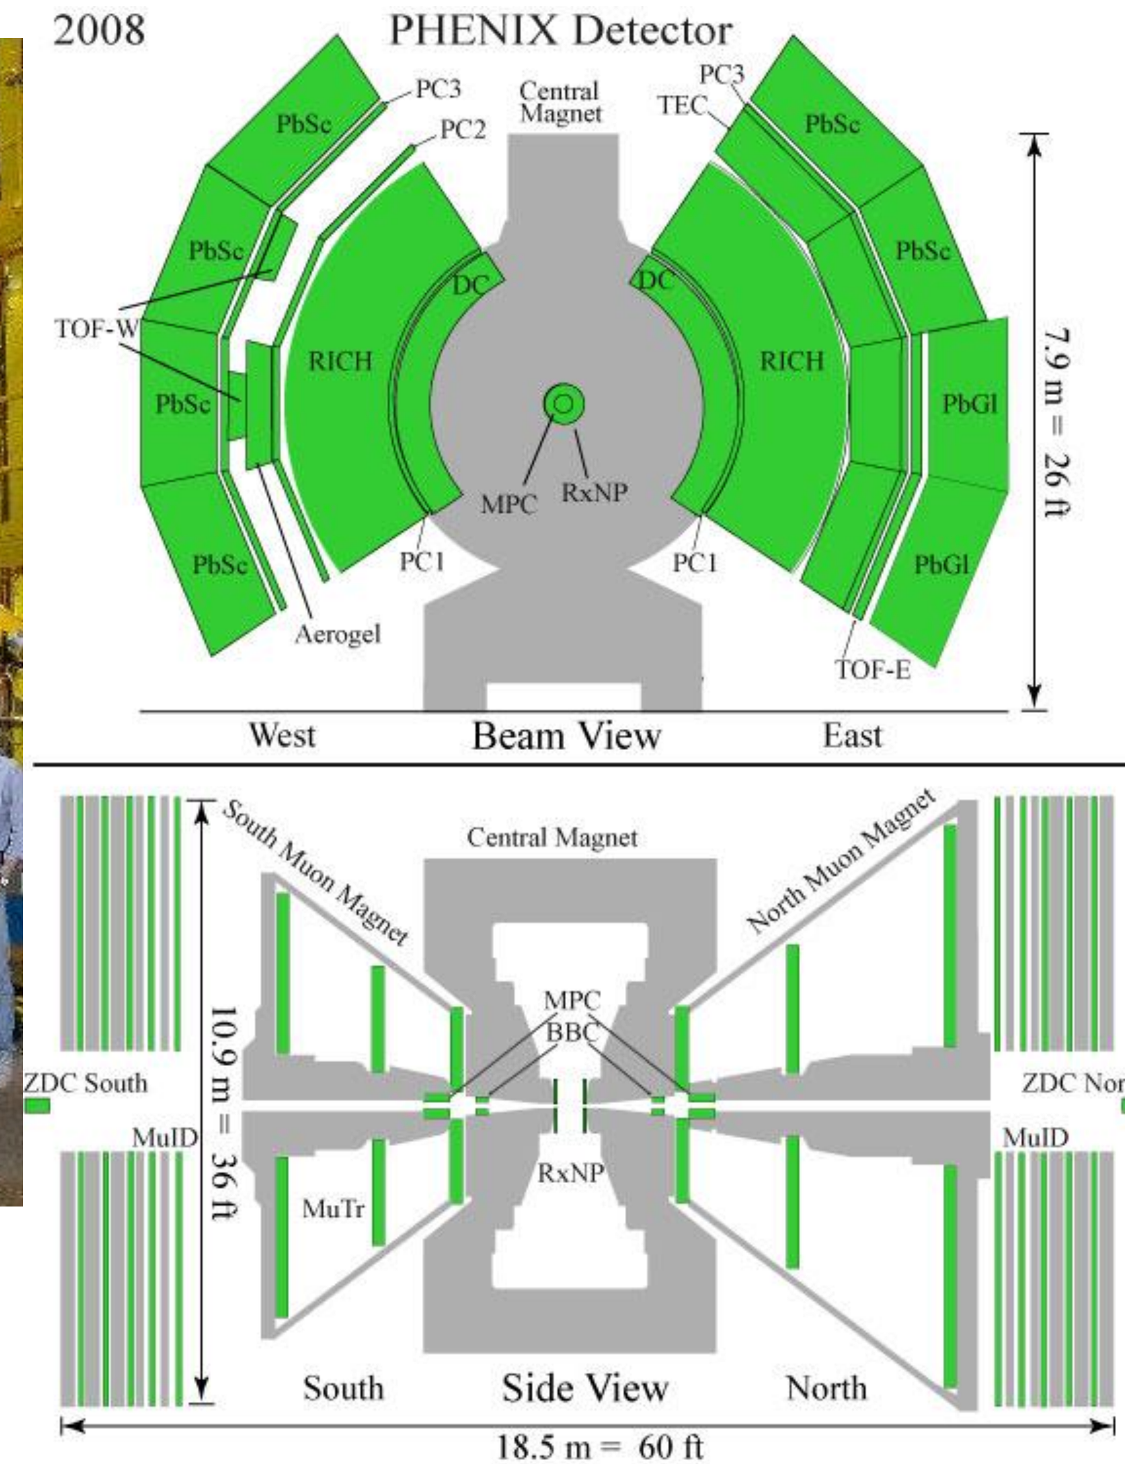

Quite a number of different detectors
4 Arms and central detectors
No hermetic coverage

PHENIX -> sPHENIX

PHENIX has been around since the early 90's

Started taking data in 2000, first Physics Run in 2001 – Run 16 is ongoing

400+ people collaboration, 200-ish publications, some "famous" ones among them

Still, the experiment is showing its age – much older than the cars most of us drive

The missing hermetic coverage is an issue for today's analyses

Limitations of the 90's designs look strange to us today

A wealth of new insights guides us what to look for next, how to build the next-gen detector

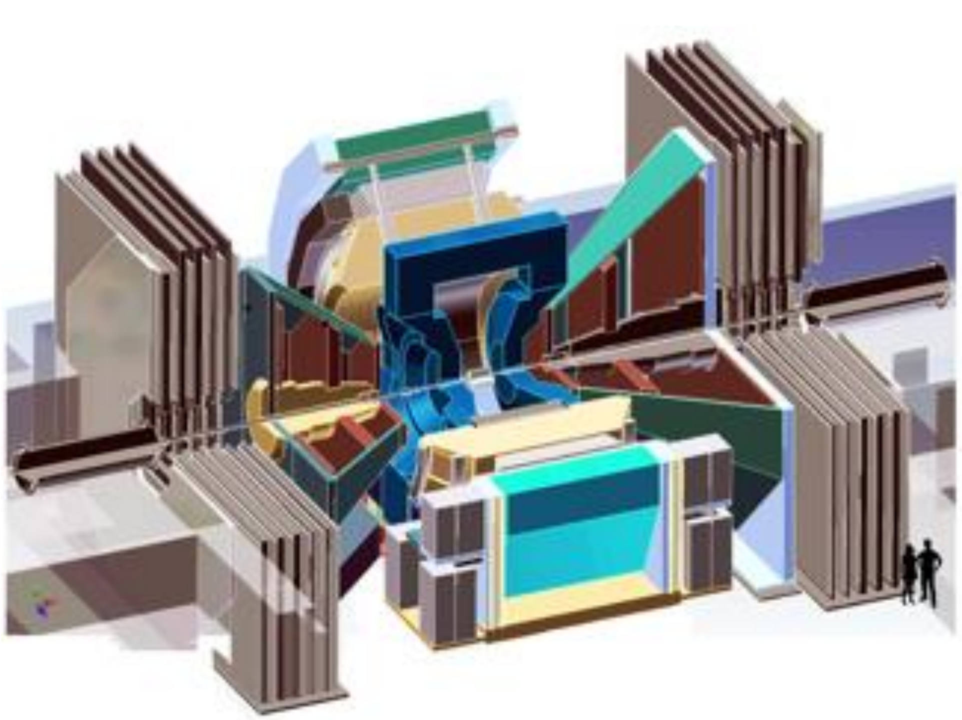

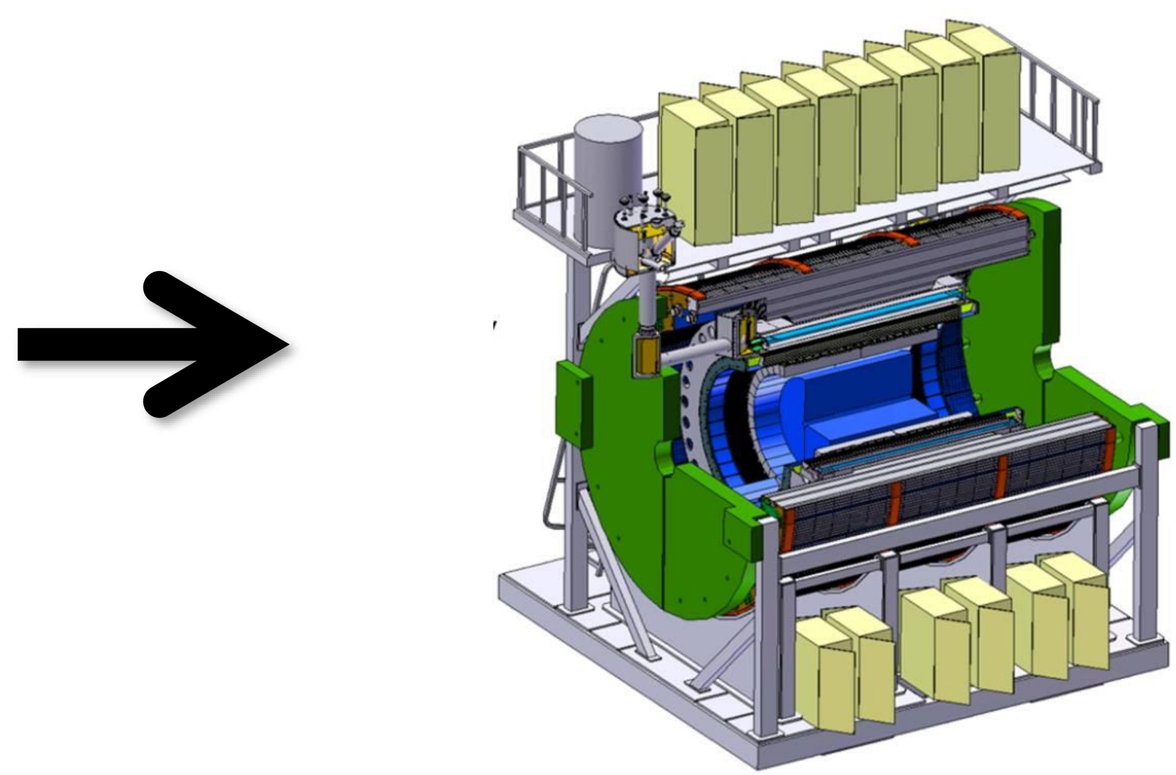

RHIC and (s)PHENIX

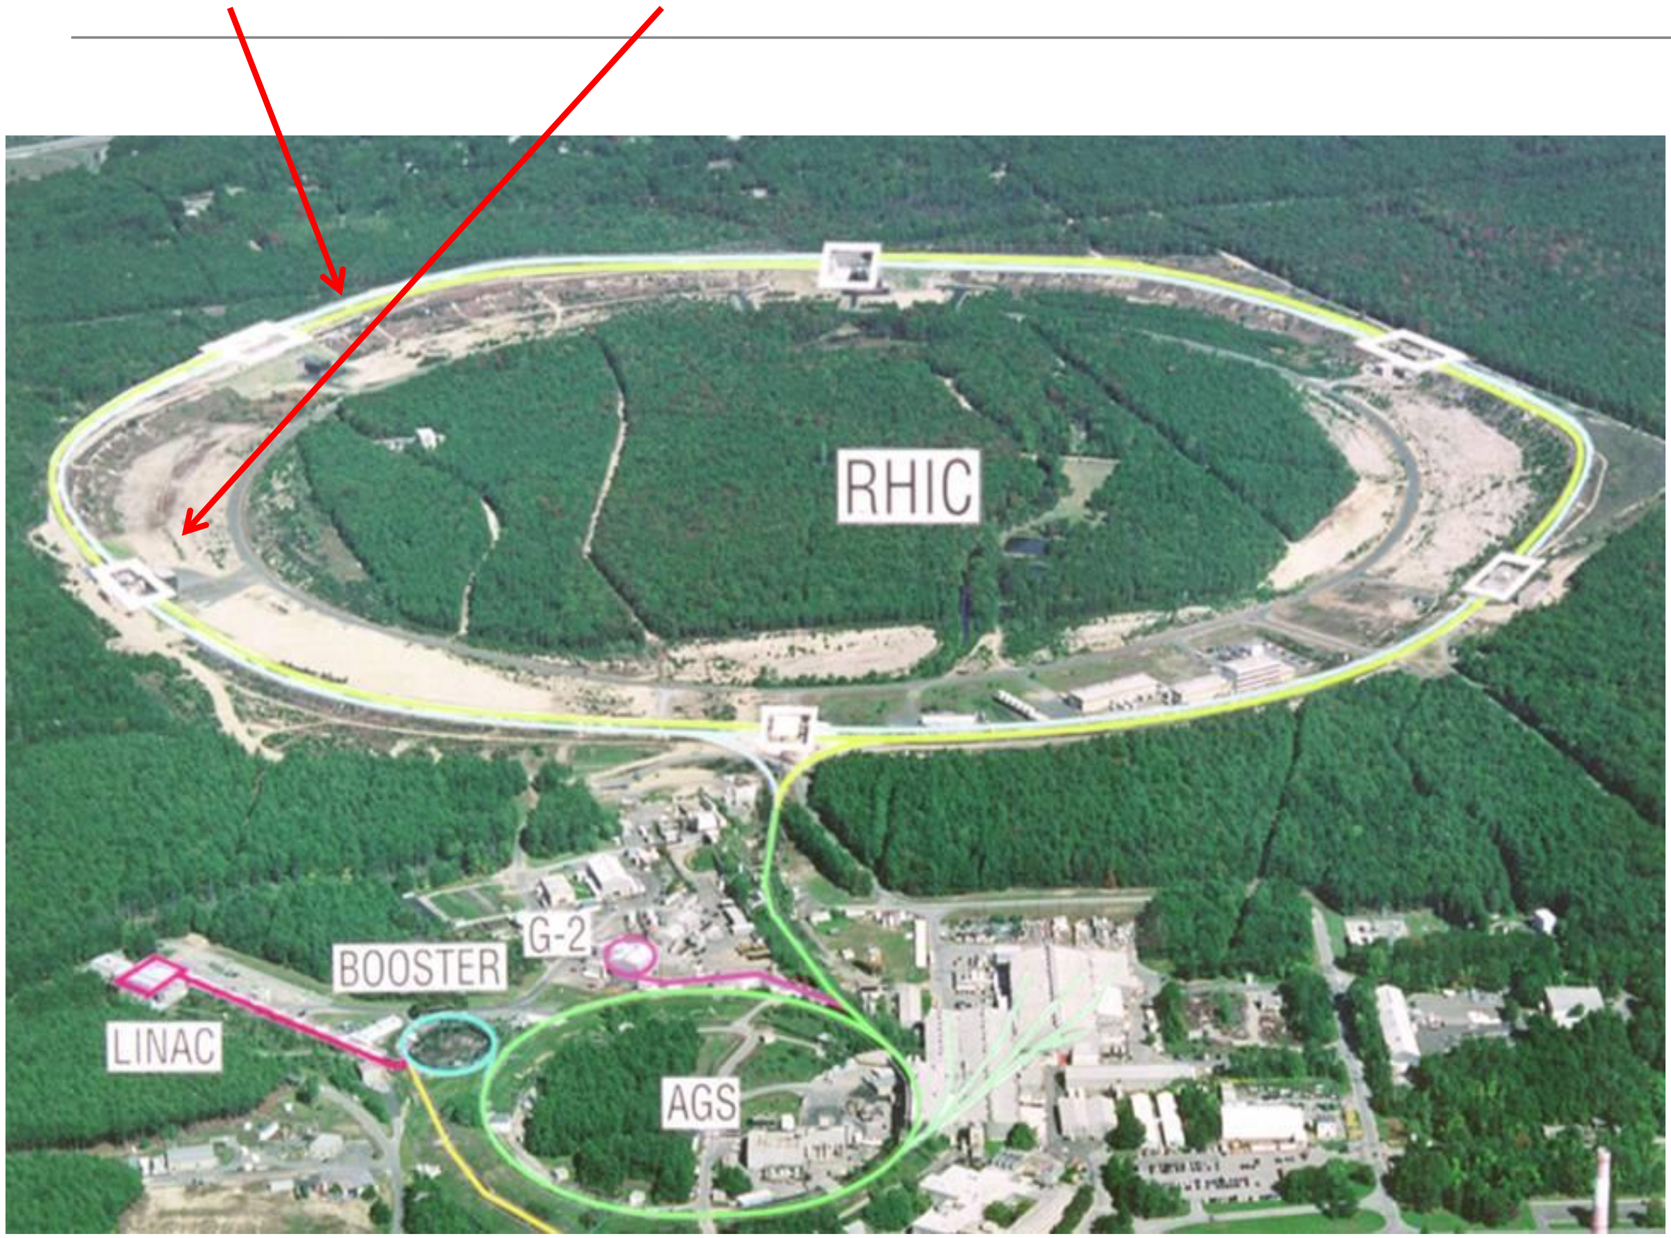

The Relativistic ***Heavy Ion*** Collider

Huge variety of ions possible - Au, Cu, ^{3}He, ^{238}U so far, but pretty much anything is possible

Polarized protons – a unique facility

500GeV/proton -> 200GeV/N for Au

Dedicated HI and polarized proton facility

PHENIX – high-rate heavy-ion experiment

Electromagnetic probes - Photons, electrons, muons

Heavy quarks & quarkonia

sPHENIX from a Physics Perspective… Jets!

- The days of measuring single particles, in however a comprehensive way, are largely over
- Most analyses look for some form of correlations between particles in individual events
- Reaction planes, the now-famous long-range correlation ”ridge”, jets
- Jets in particular is what carries the physics these days – probing the medium

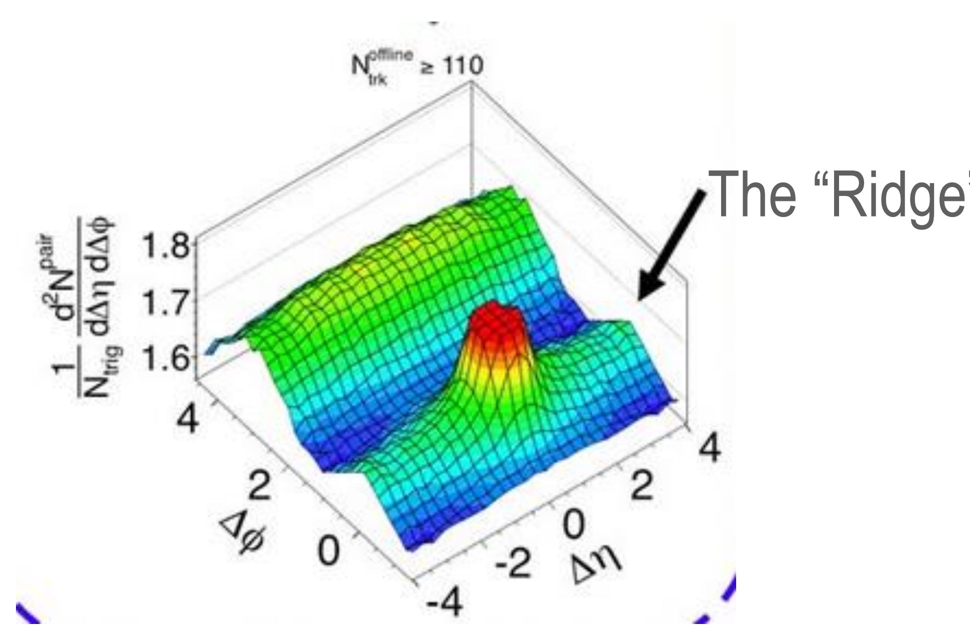

Jets

Jets are made in hard QCD scattering of quarks or gluons

You get an initial pair of quarks with (mostly) opposite momentum

Quarks hadronize – they produce a “spray” of particles in narrow cone, the “jets”

So we are looking for those back-to-back cones of particles

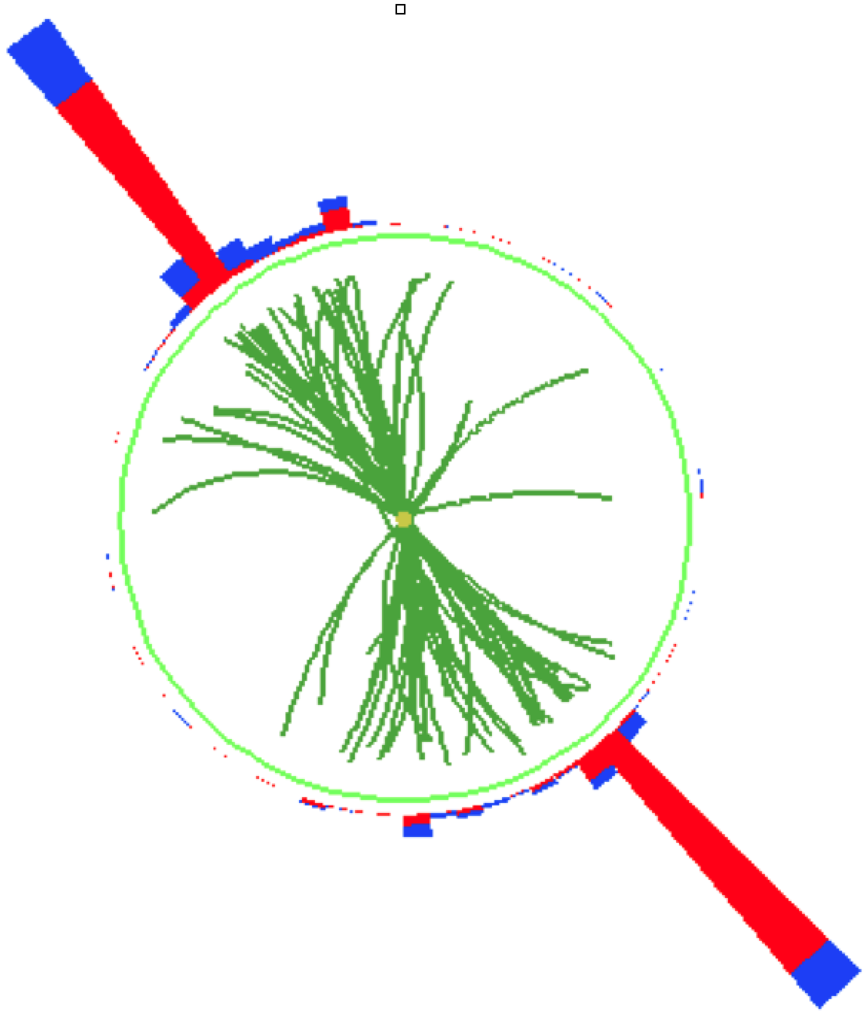

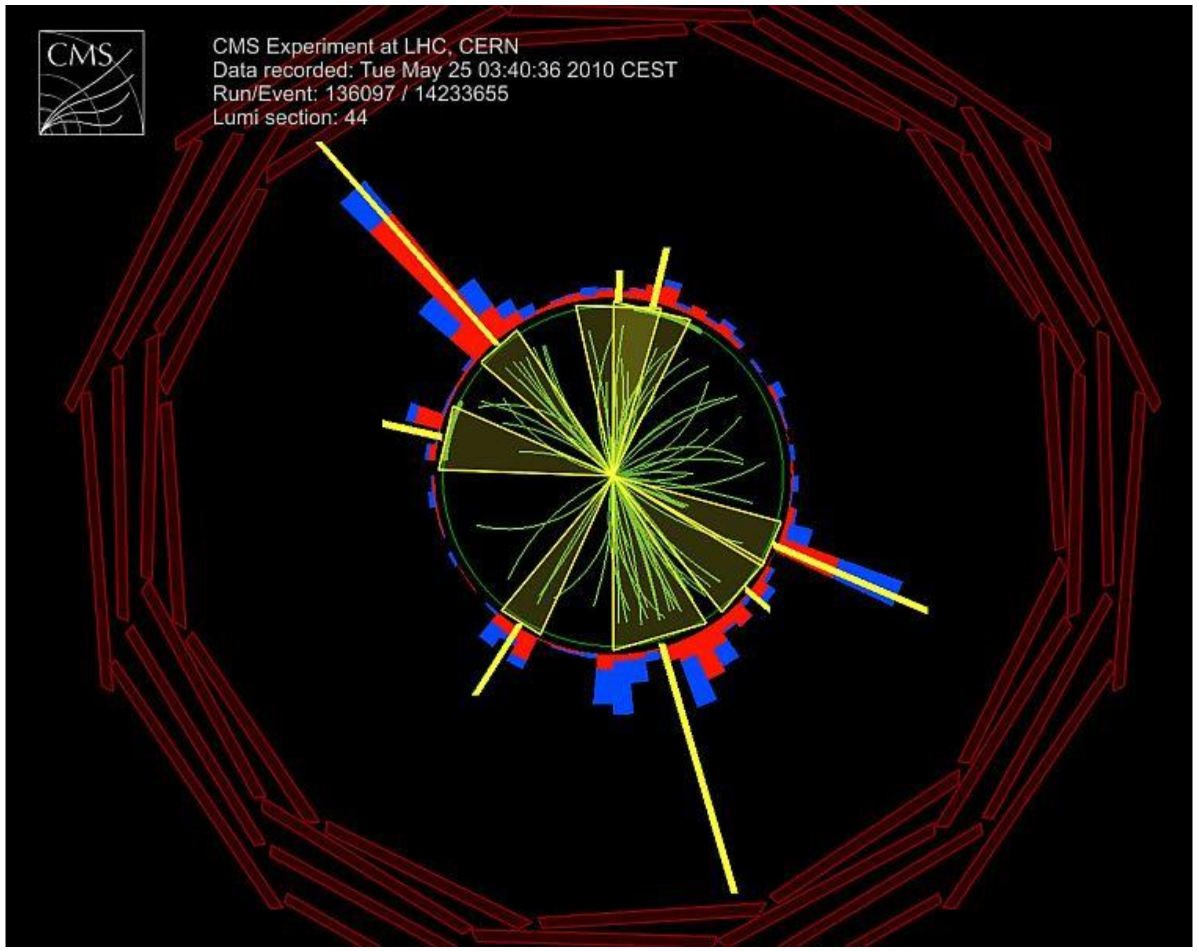

So, what do we want?

While we have a number of jet measurements, the PHENIX acceptance is too small (and has too many holes) to be a good jet detector

So what we want is:

- full azimuthal coverage, large coverage in rapidity (+- one unit)
- decent tracking
- electromagnetic and hadronic calorimetry
- compact

So let’s go and upgrade the experiment!

Why didn't we build PHENIX like that in the first place?

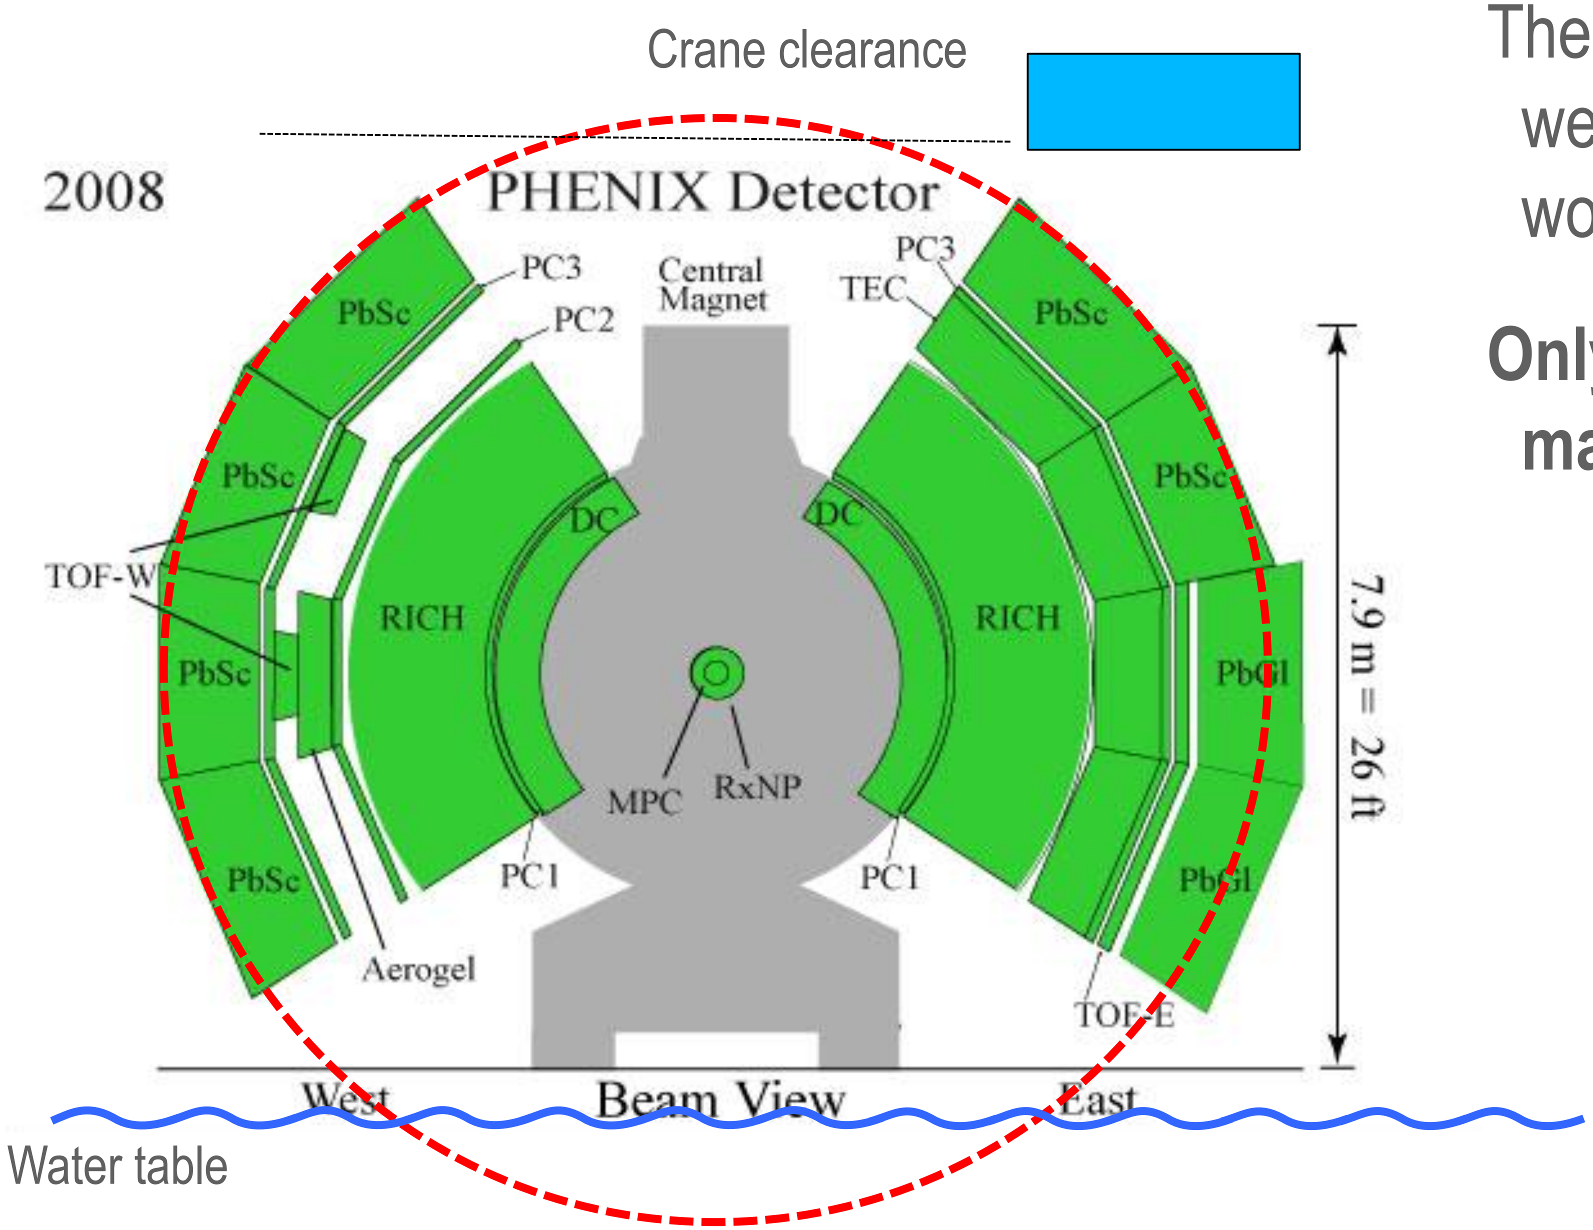

The available materials at the time were such that a full-azimuth detector would not fit into the experimental hall

Only new techniques and materials make this possible, 20+ years later

"Lake PHENIX" level

If one digs just a few inches, there's water…

sPHENIX – the Concept

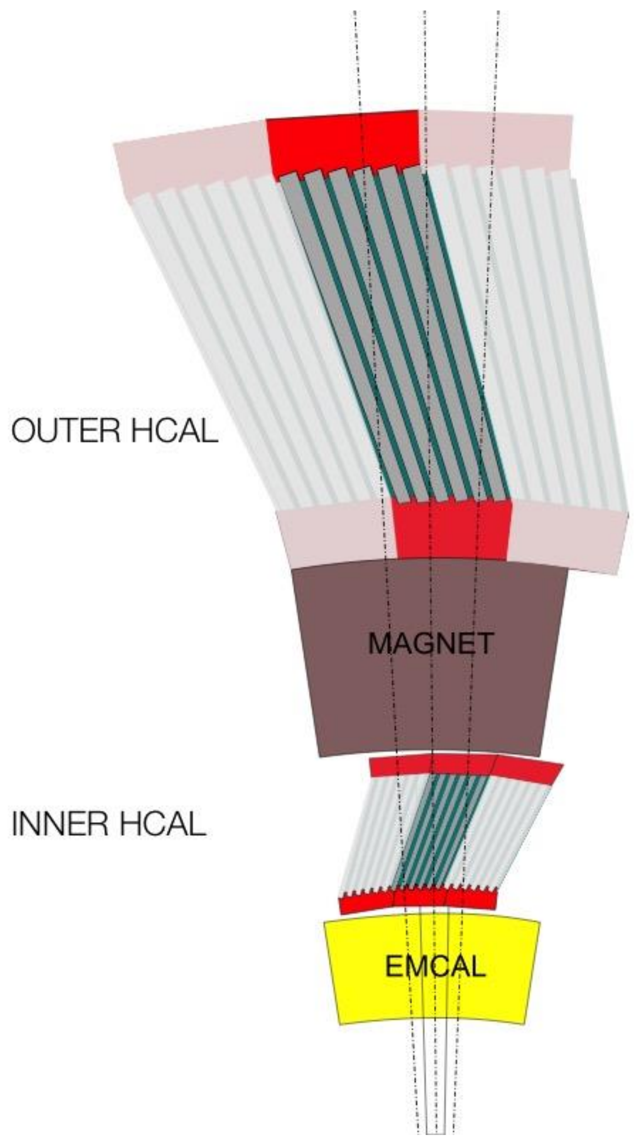

- Outer HCAL $\approx 3.5\lambda_I$
- Magnet $\approx 1.4X_0$
- Inner HCAL $\approx 1\lambda_I$
- EMCAL $\approx 18X_0 \approx 1\lambda_I$
- Tracker

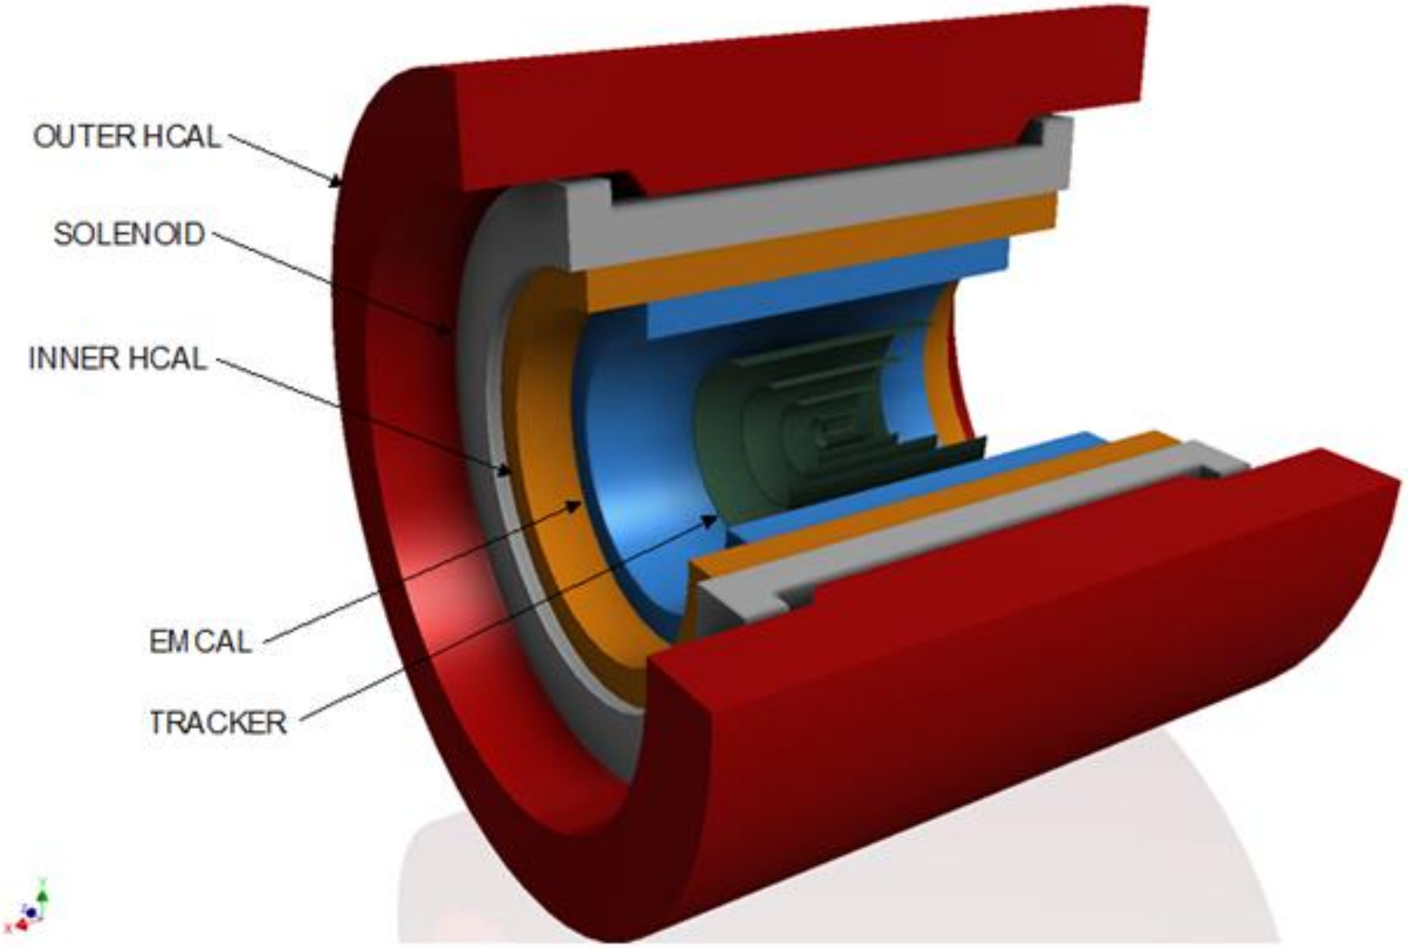

HCAL steel and scintillating tiles with wavelength shifting fiber

2 longitudinal segments.

An Inner HCal inside the solenoid.

An Outer HCal outside the solenoid.

$\Delta\eta \times \Delta\varphi \approx 0.1 \times 0.1$

2 x 24 x 64 readout channels

$\sigma_E/E < 100\%/\sqrt{E}$ (single particle)

SiPM Readout

The BaBar Magnet

The EmCal (this is is really the special item…)

A lot of the space savings (compared to PHENIX) are found in the Emcal

In the 90's, the standard EmCal was Lead Scintillator or Lead glass

Comparable radiation lengths and Moliere radii for both – **28mm and 37mm**

Metallic tungsten would be cool, **X_0= 3.5mm**!

But impossible to machine, melt, form…

Then a solution came from the most unlikely place – a golf club manufacturer

Instead of using metallic tungsten, use a paste of tungsten powder and epoxy

Not as dense as pure tungsten, but good enough – **X_0 ~ 7mm**

Can be shaped in any form (e.g. golf club heads), is cheap (uses "scrap" tungsten)

So we build the EmCal out of that!

Tungsten Powder EmCal Modules

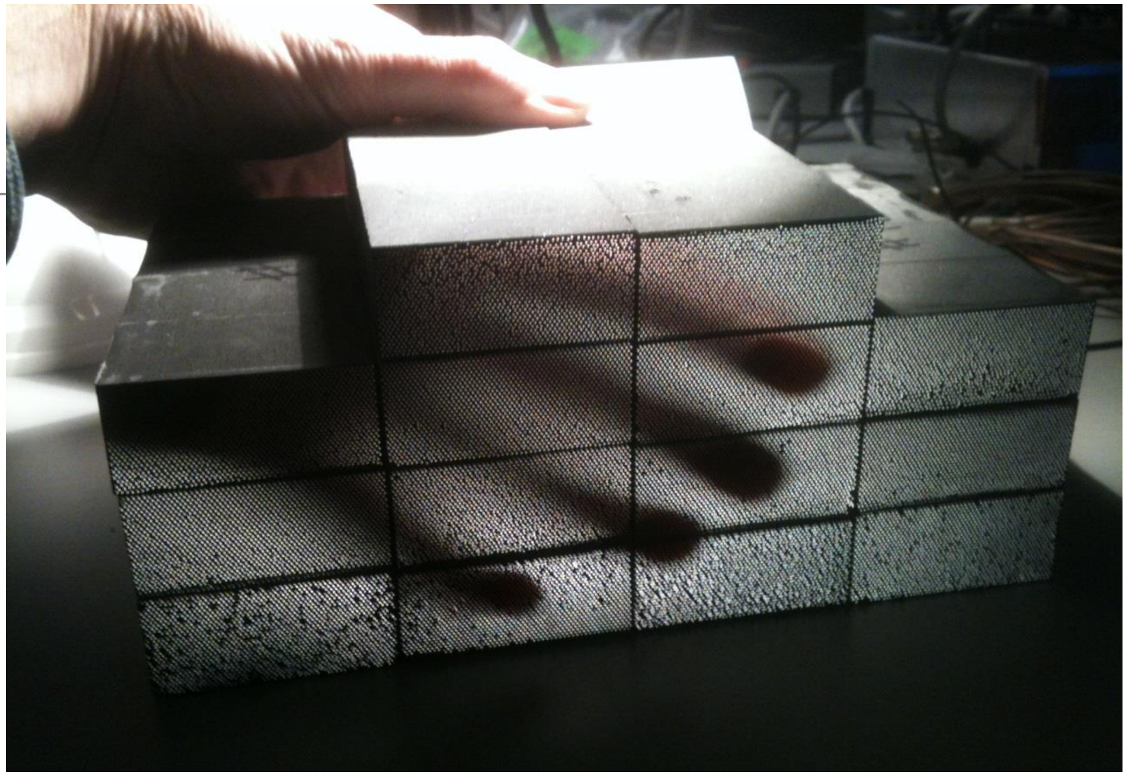

Heavier than solid Lead

We are producing Emcal modules in-house (UIUC), and at a company called THP

Avg density ~ 9.3 g/cm^3 ± ~ 0.1 g/cm^3

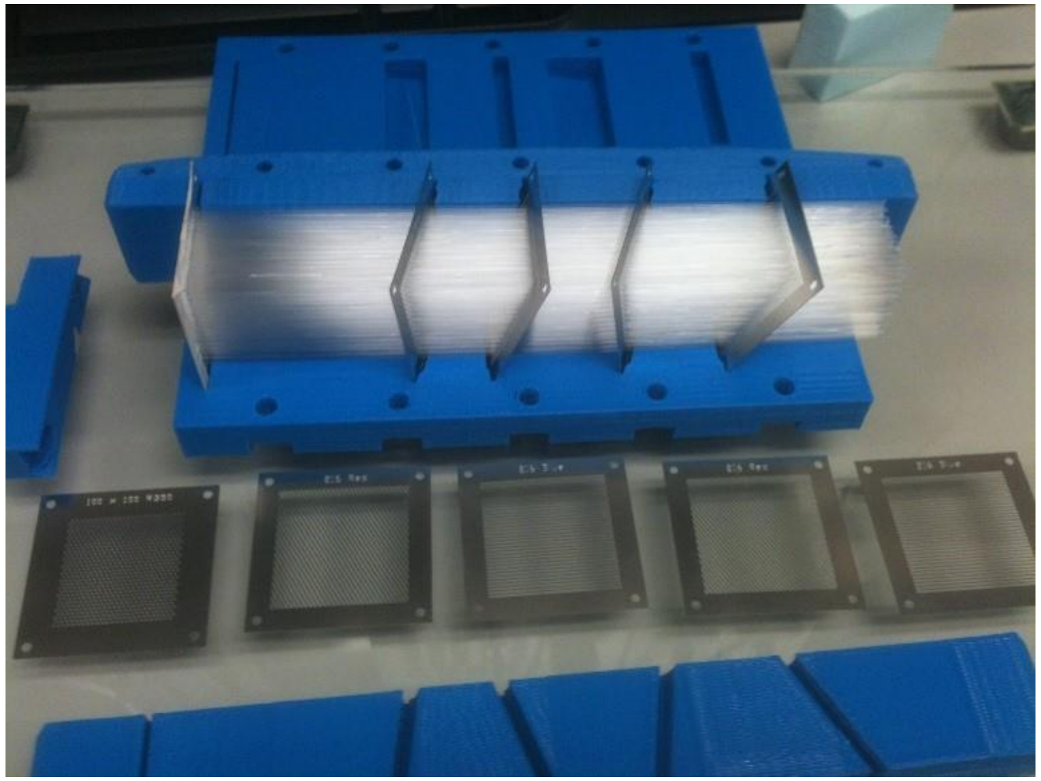

The mold with scintillating fibers

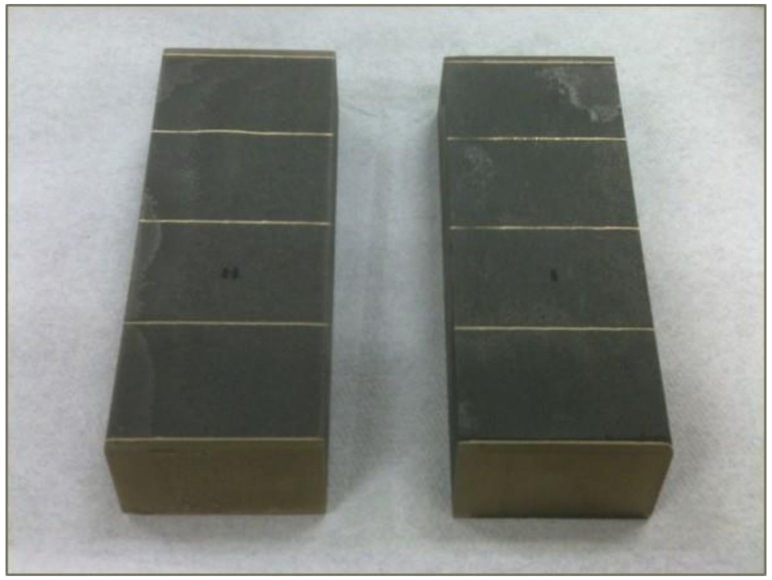

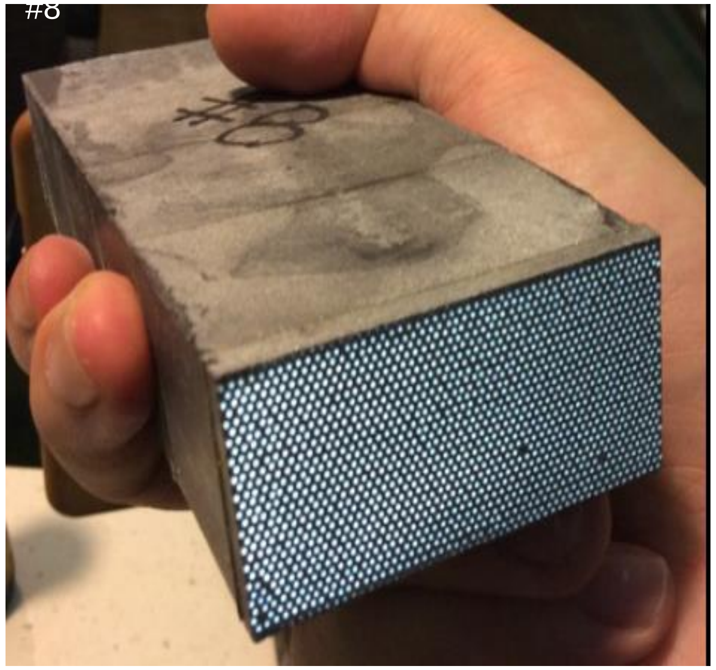

All calorimeters are read out with SiPMs

Talking about the Readout of all that…

- PHENIX has never been afraid of high data rates and volumes
- I personally claim a portion of that fame…
- We fully intend not to sacrifice data and "take all we can"
- Current design goal is 15KHz of event rate
- This is possible right now, but will be less of an issue in 2020+

sPHENIX will have an estimated 15-25 GByte/s fully compressed data stream

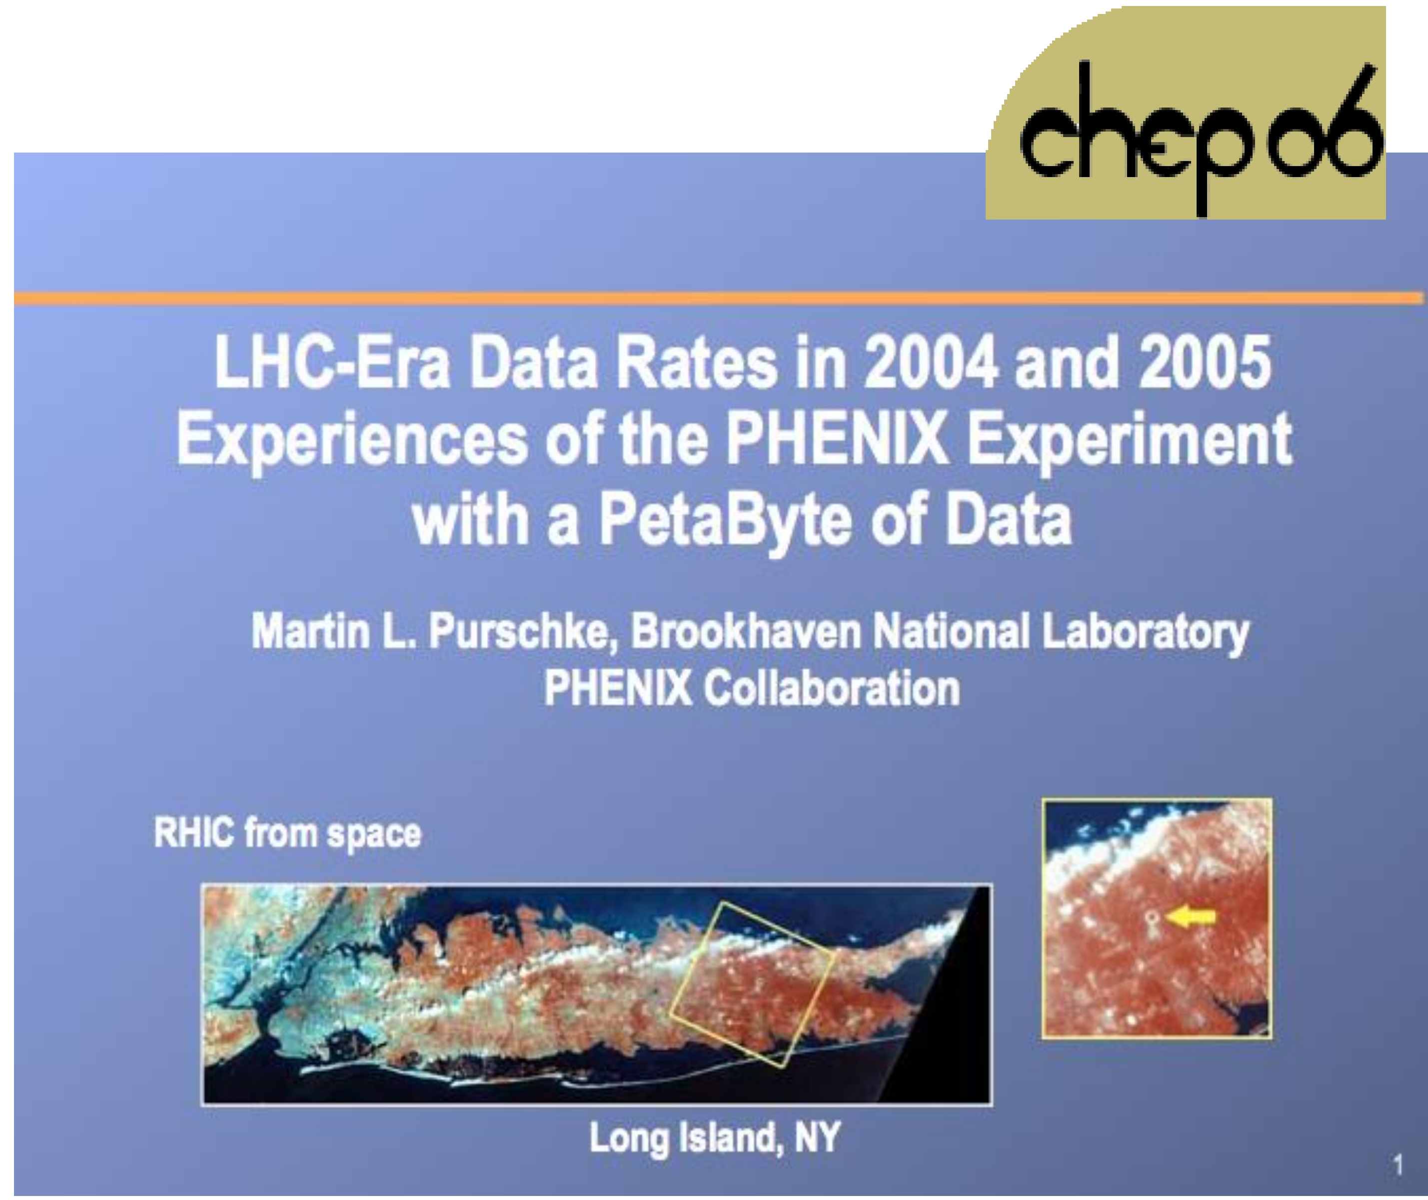

My opening slide from the CHEP in 2006

The sPHENIX DAQ

- As much as possible, we will re-use the existing PHENIX DAQ
- We have been asked by review committee members why we wouldn't want to upgrade a 15yr old DAQ
- Answer: We already did that, we have just the system that we would build today
- At least for the front-end!
- Legacy gear from older detectors goes out the window
- Working and very modern front-end, lots of experience
- We have traditionally taken the highest data rates in the field
- 15KHz event rate envisioned, 14 have been demonstrated
- Most of the work is in the back-end, event builder, etc
- Test beam setup used an improved DAQ system already

Much less variety

- PHENIX had 14 different detector system, each with its assorted front-end electronics
- Herculean integration efforts early on
- sPHENIX, by conscious design, has only a few and common technologies
- Concentrate on uptime

DAQ Overview

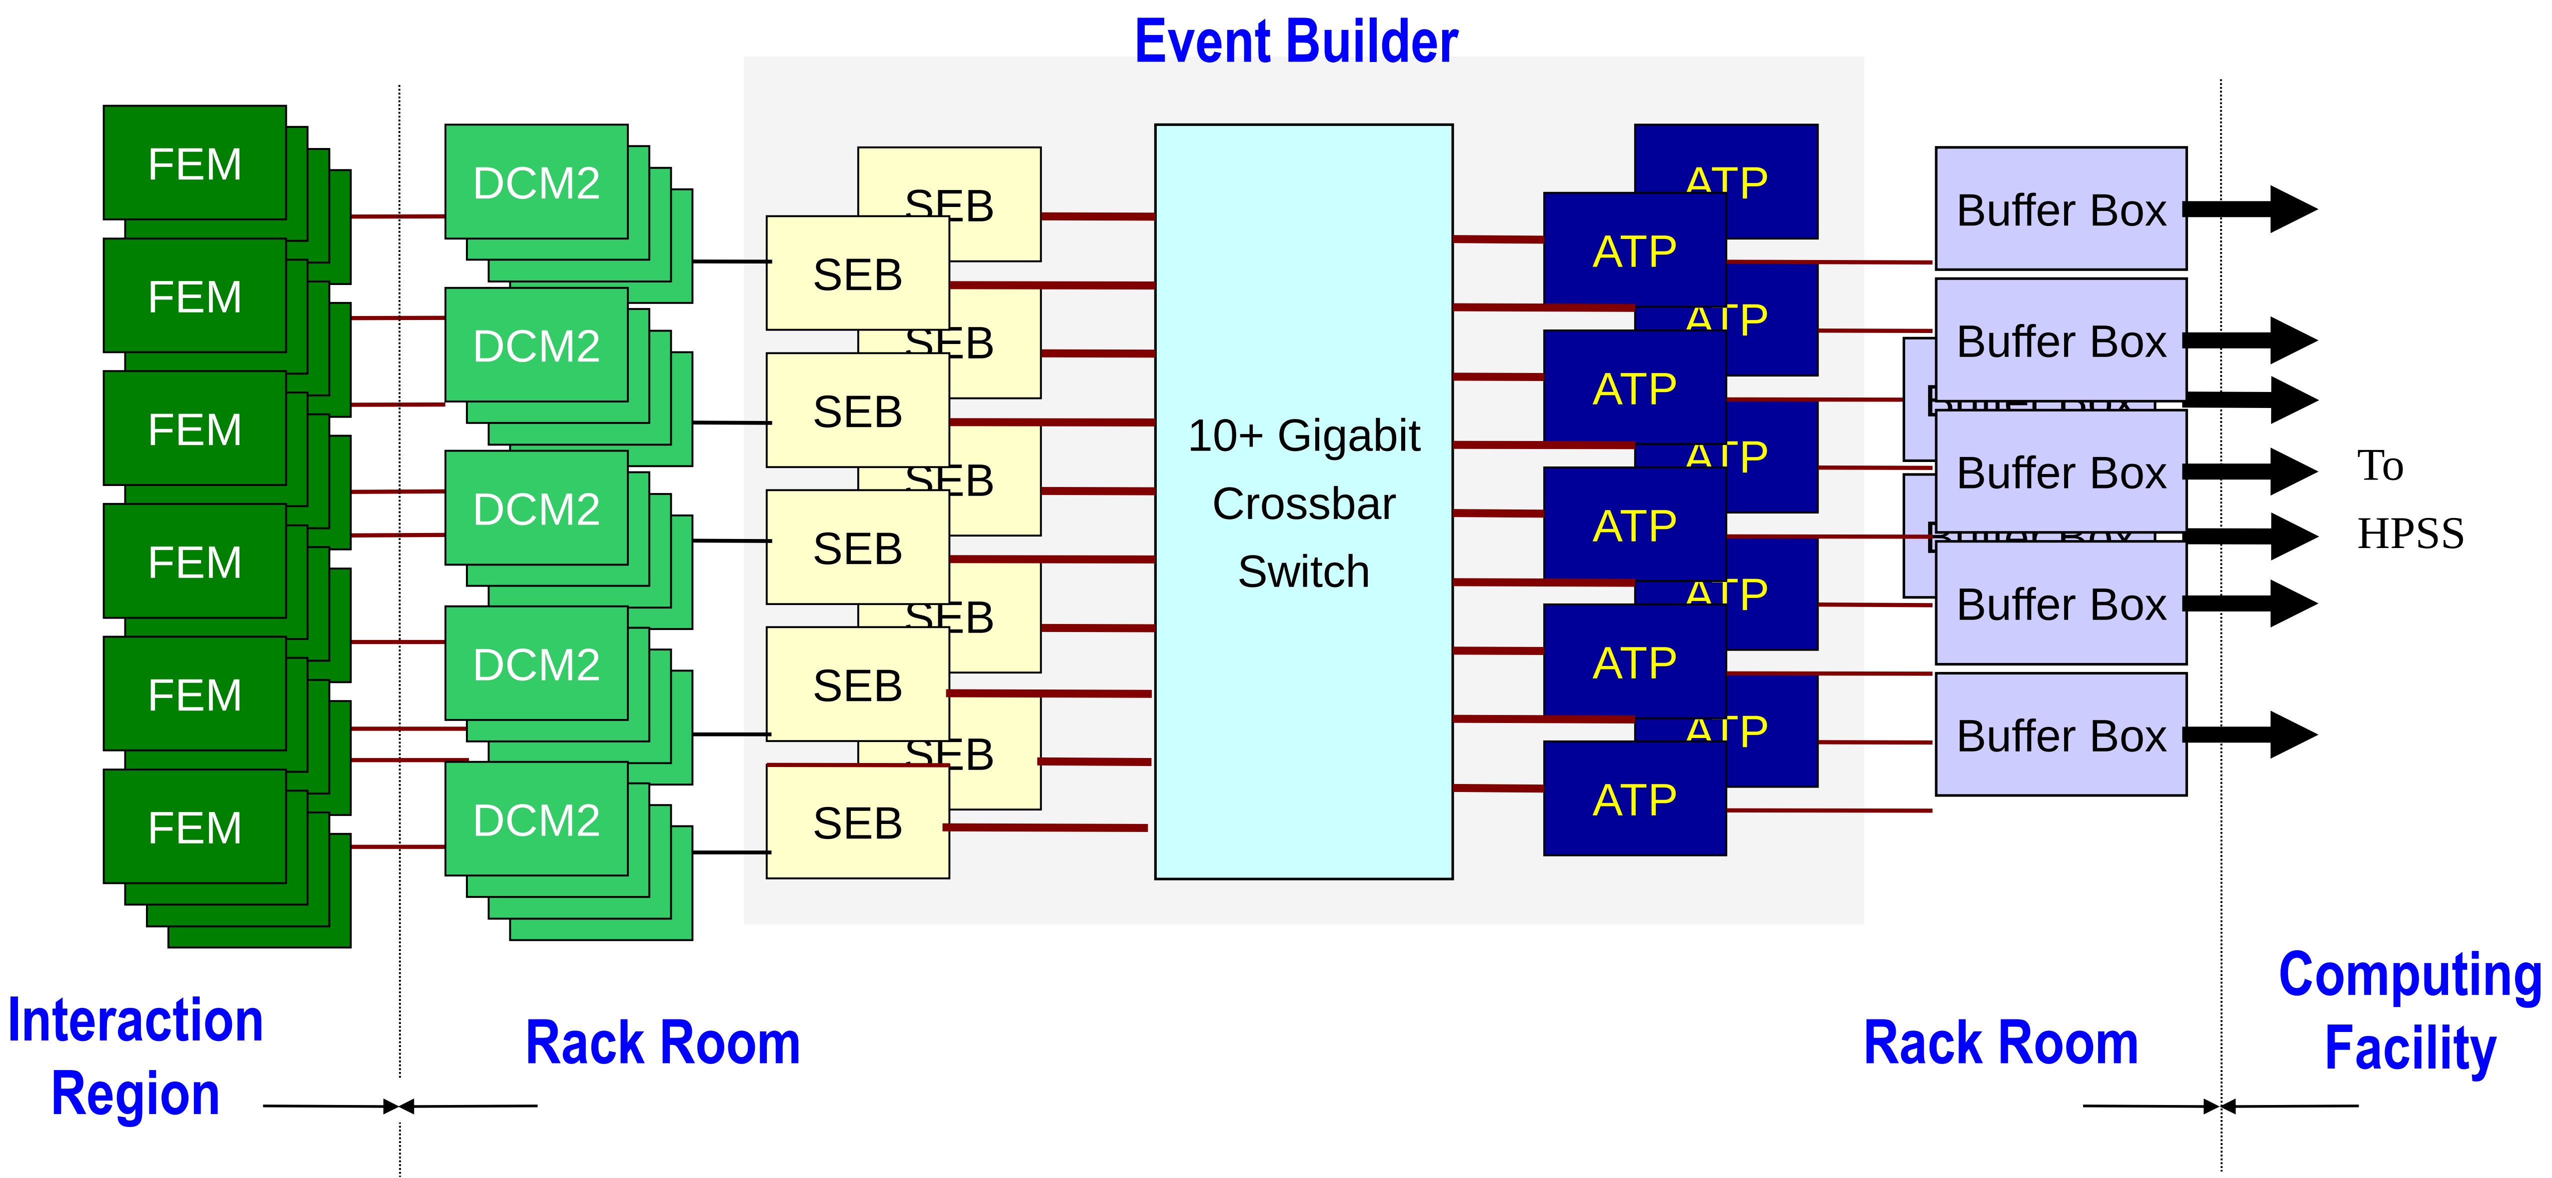

- DCM-2 receives data from digitizer, zero-suppresses and packages
- SEB collects data from a DCM group (~35)
- ATP Assembles events and compresses data (~60)
- Buffer Box data interim storage before sending to the computing center (7)

Slot Controllers

- Provide monitoring and control for EmCal and HCal Interfaces.
- Temperature and leakage stabilizer functions.
- Temperature gain correction critical for SiPMs
- Near-prodution boards used for EmCal and Hcal during beam test.

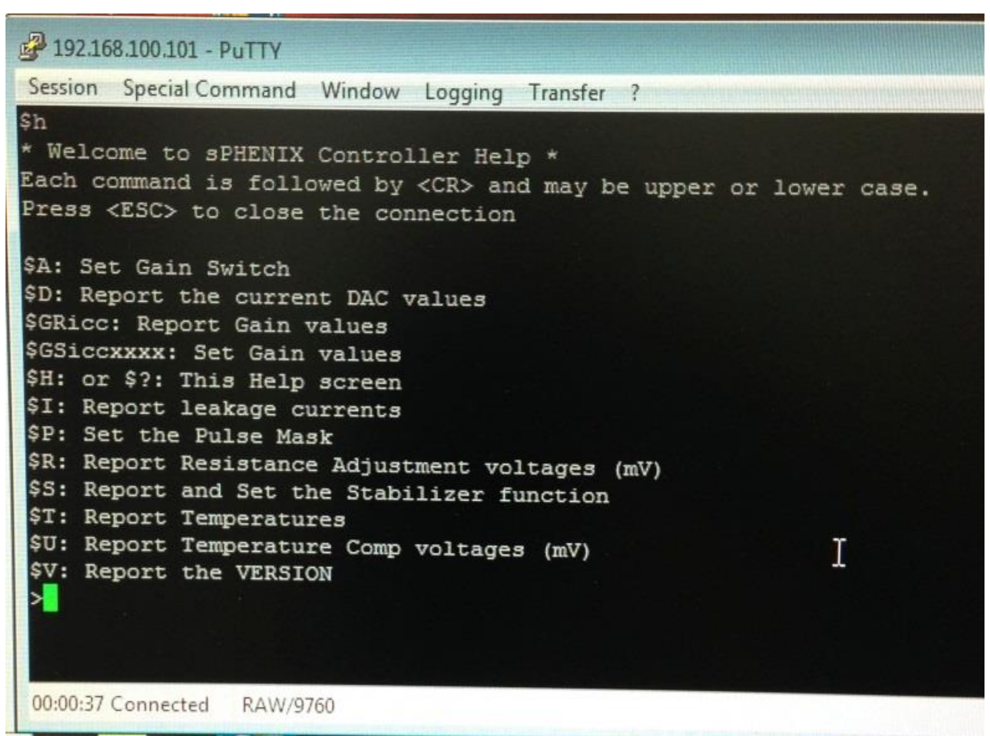

Digitizer/ADC Boards, XMIT, Controller

- Goal is to run the next test beam (January 2017) with new digitizer electronics
- 64 channels per board
- 14 bit ADC
- 60 MHz sampling frequency
- XMIT board interfaces 4 ADC boards to DCM-II
- Controller board interfaces to SEB

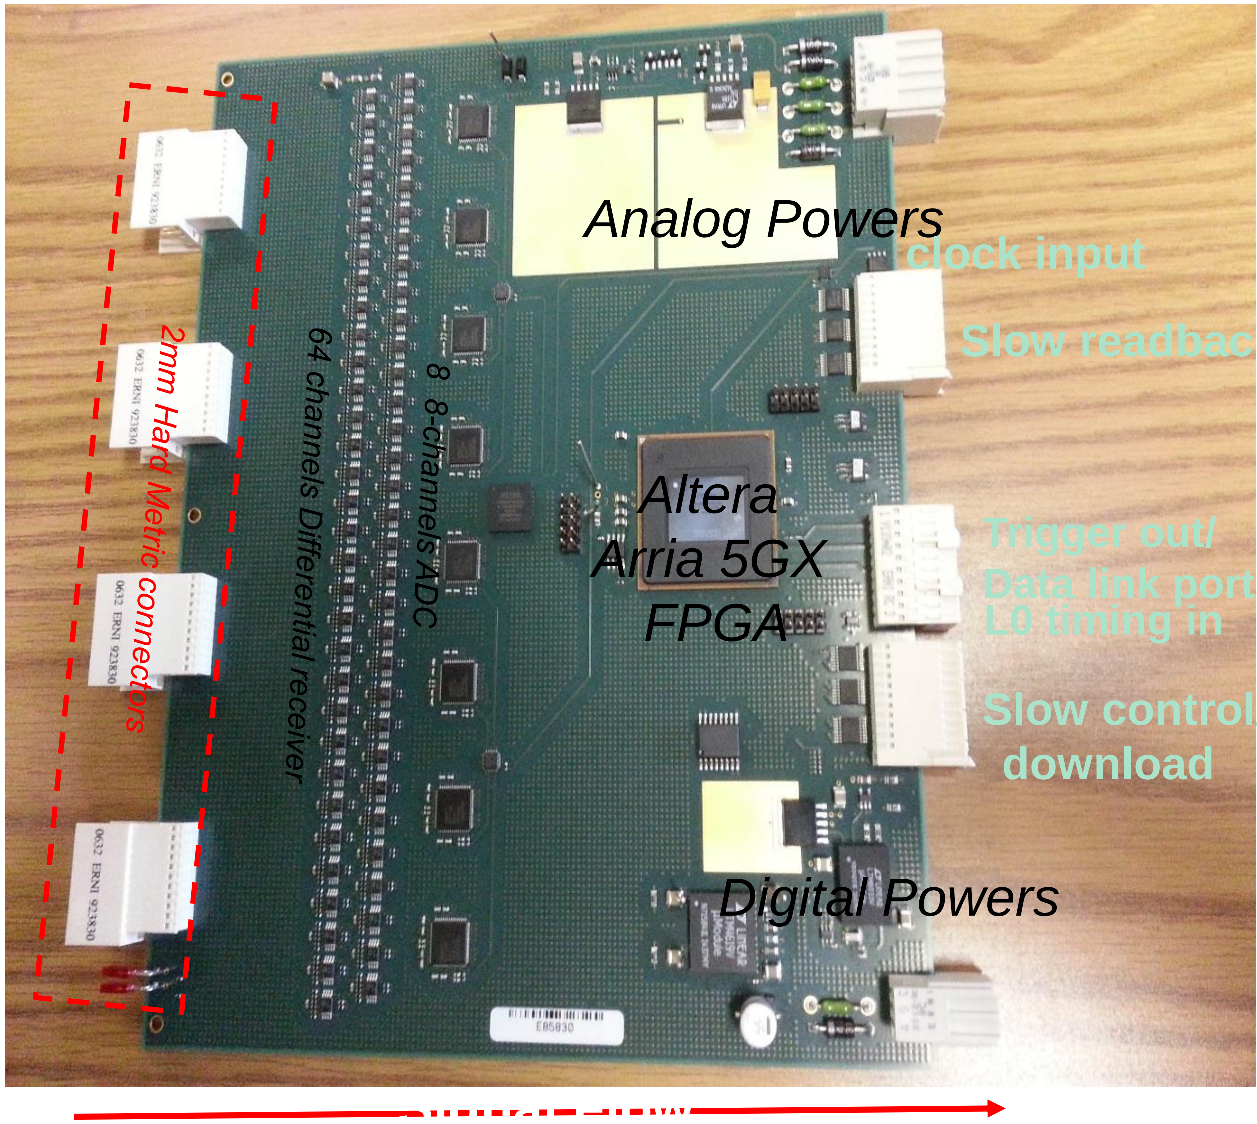

DAQ Components – DCM2

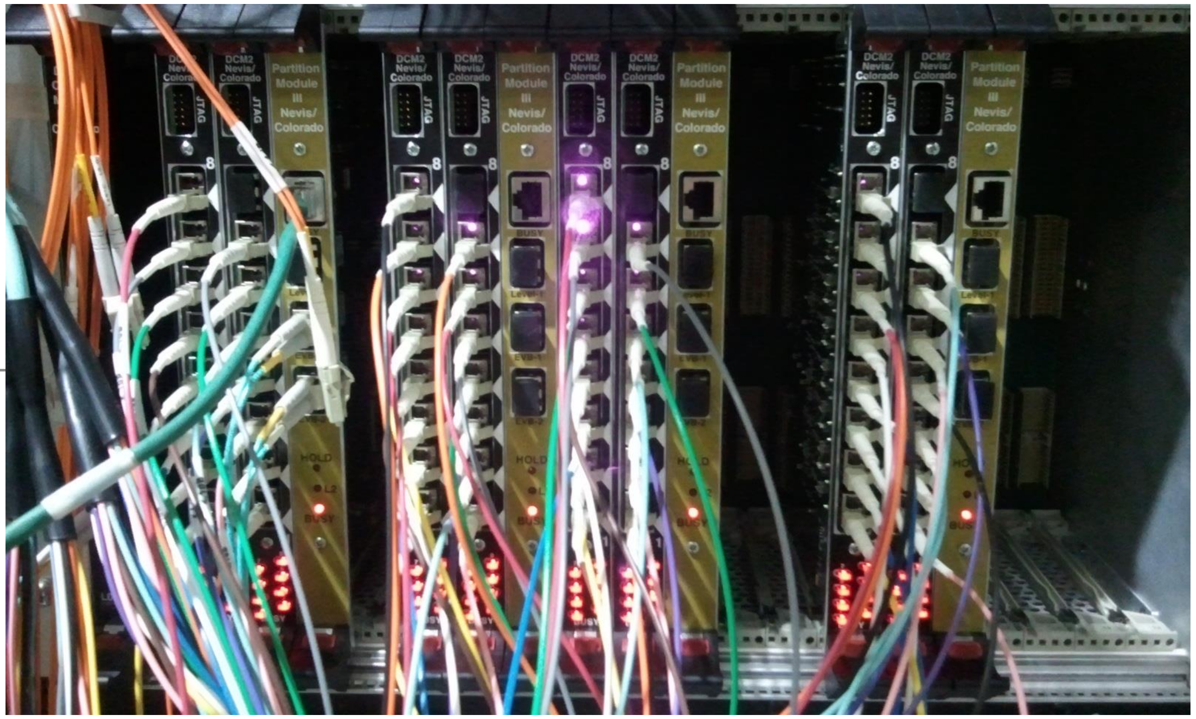

- The Data Collection Module V2 is a modern, mature board commissioned in PHENIX Runs 14,15
- "just as we would have built it" for sPHENIX
- FPGA code, tools, existing configurations to draw from
- Configuration tools / description language in place
- O(200) units
- Production-grade boards in hand for test beams, R&D, etc as soon as the PHENIX Run ends (end of June)

FNAL Testbeam Impressions

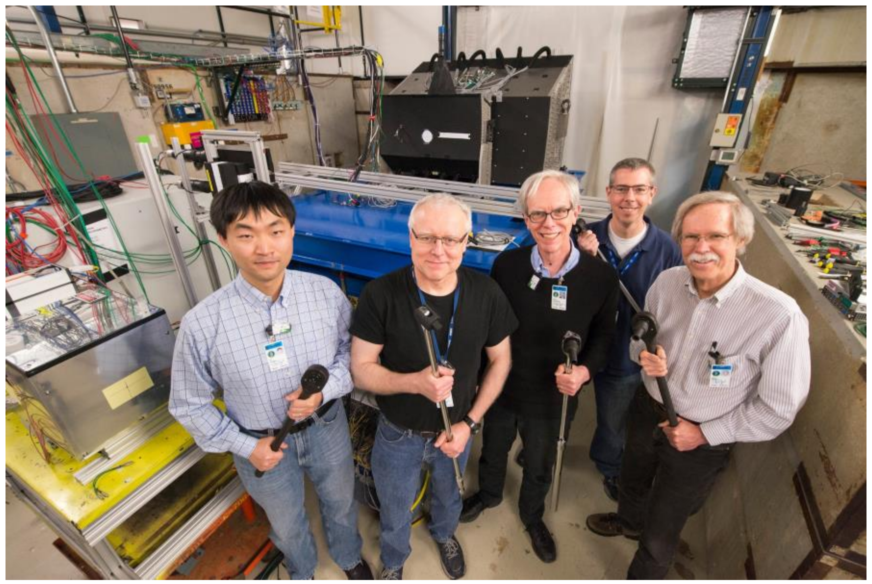

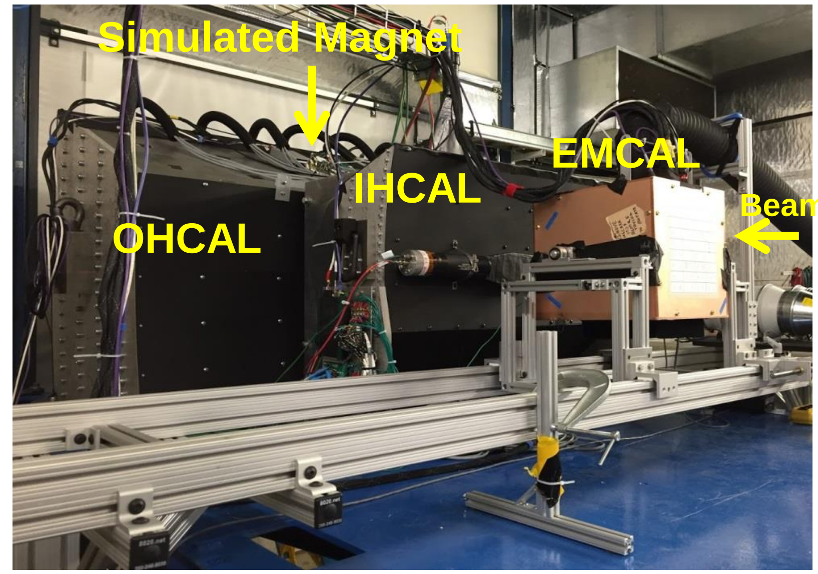

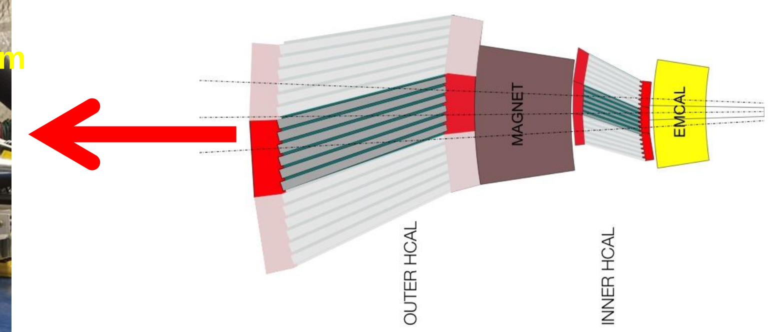

- Need to verify our simulation results with test beams before production
- Had a really good run with few problems
- Unified data taking for all (or at least most) projects with RCDAQ
- Rcdaq is in use in sPHENIX and the EIC orbit (BNL, UIUC, SBU, Yale, GSU, even ATLAS ZDC test beam/calibration)
- Produces the eventual sPHENIX data format (PRDF)
- Ideal to get students to be proficient working with the raw data

Metadata Logging

Martin says: *"Log everything you can think of. And then log some more."*

- Log automatically. Do not rely on human input.
  How was the DAQ configured? What was the HV? Were the lights on? What was the ambient temperature? Was something upstream in the beamline?
- Many things just for "forensics" purposes in case something doesn't make sense
- I truly believe in taking pictures – a pic captures *everything*
- Metadata captured in a database, but also in the raw data
- Info cannot get lost, easy to access if you have the raw data file

Metadata Logging Example – FTBF beam params

We had a RPi that was aware of the spill end – we fed it a copy of the spill signal into a high-tech board

Energy	S:MTNRG	=	-16	GeV
	F:MT6SC1	=	10129	Cnts
Beamline Counters	F:MT6SC2	=	9013	Cnts
	F:MT6SC3	=	7829	Cnts
	F:MT6SC4	=	0	Cnts
Beamline Counters	F:MT6SC5	=	158725	Cnts
Emcal table h/v position	E:2CH	=	248.3	mm
Emcal table h/v position	E:2CV	=	11.33	mm
Ambient temp & humidity	E:2CMT6T	=	73.44	F
Ambient temp & humidity	E:2CMT6H	=	39.7	%Hum
Cher. pressures	F:MT5CP2	=	1.242	Psia
Cher. pressures	F:MT6CP2	=	1.492	Psia

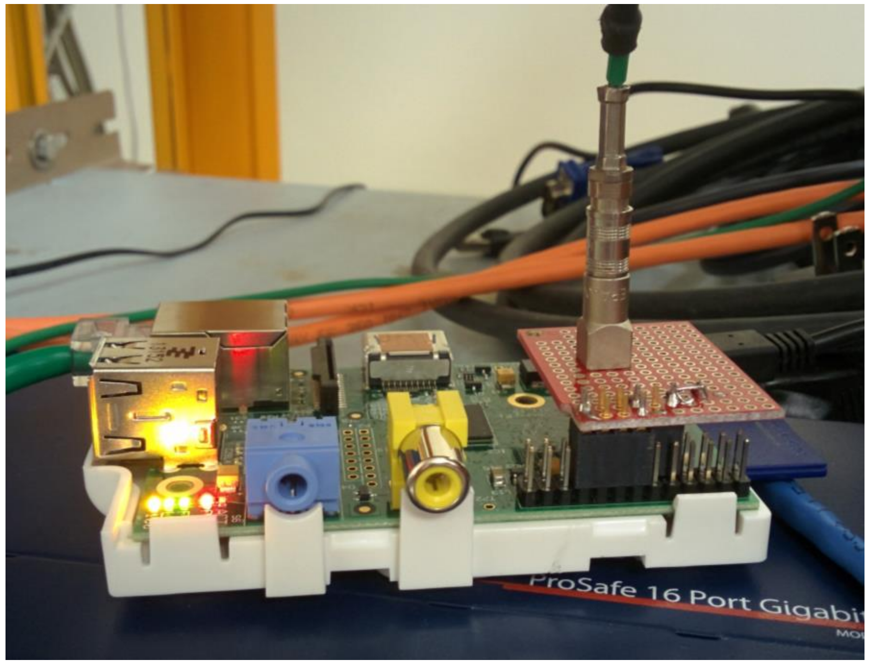

Very useful for “forensics”

“It appears that the distributions change for Cherenkov1 at 1,8,12,and 16 GeV compared to the other energies. It seems that the pressure is changed. […] Any help on understanding this would be appreciated.”

Martin: “Look at the info in the data files”:

```
$ ddump -t 9 -p 910 beam_00002298-0000.prdf
S:MTNRG   = -1      GeV
F:MT6SC1 =  5790        Cnts
F:MT6SC2 =  3533        Cnts
F:MT6SC3 =  1780        Cnts
F:MT6SC4 =  0           Cnts
F:MT6SC5 =  73316       Cnts
E:2CH     =  1058   mm
E:2CV     =  133.1  mm
E:2CMT6T =  73.84  F
E:2CMT6H =  32.86  %Hum
F:MT5CP2 =  .4589  Psia
F:MT6CP2 =  .6794  Psia
```

```
$ ddump -t 9 -p 910 beam_00002268-0000.prdf
S:MTNRG   = -2      GeV
F:MT6SC1 =  11846       Cnts
F:MT6SC2 =  7069        Cnts
F:MT6SC3 =  3883        Cnts
F:MT6SC4 =  0           Cnts
F:MT6SC5 =  283048      Cnts
E:2CH     =  1058   mm
E:2CV     =  133    mm
E:2CMT6T =  74.13  F
E:2CMT6H =  37.26  %Hum
F:MT5CP2 =  12.95  Psia
F:MT6CP2 =  14.03  Psia
```

More Forensics

"Was the scintillator test contraption in the beam in run 2743? There is a higher fraction of showering than before."

Look at the cam pictures we automatically captured for each run:

```
$ ddump -t 9 -p 940 beam_00002742-0000.prdf > 2742.jpg
$ ddump -t 9 -p 940 beam_00002743-0000.prdf > 2743.jpg
```

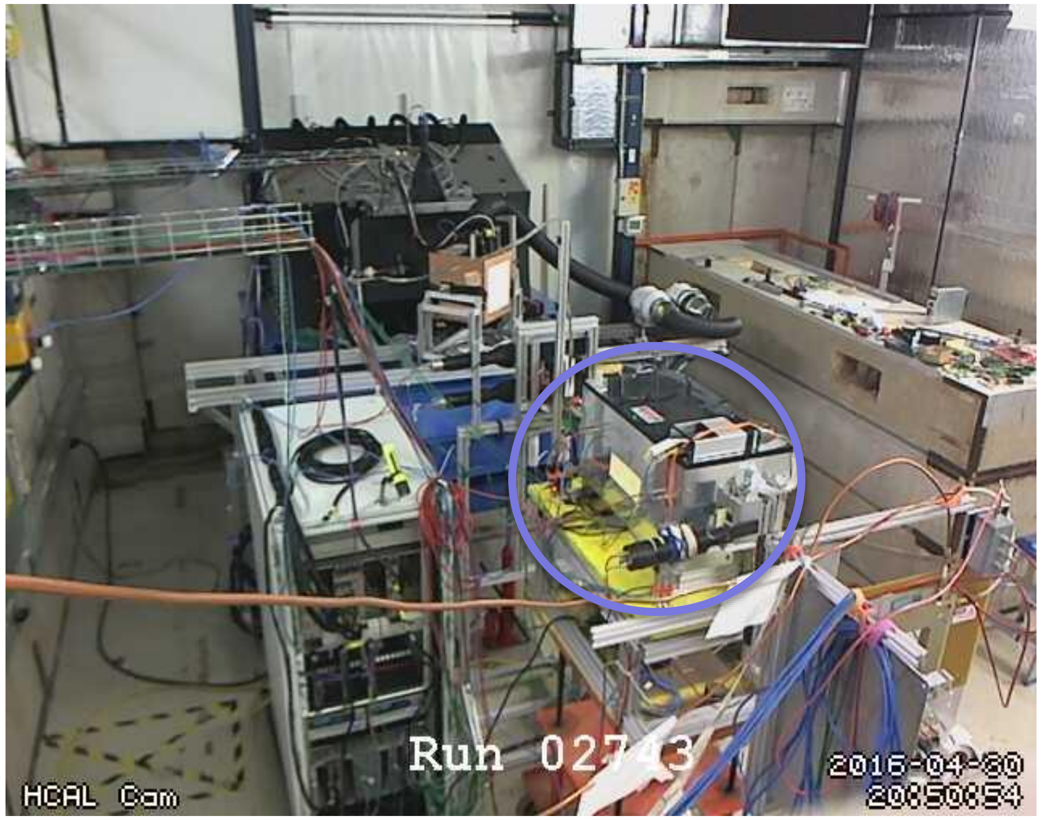

Example: Temperature compensation

The FTBF A/C was hard at work – temperature cycles at the 30 mins level

Estimated -3.6%/^{0}C gain variation in the SiPMs

Need to see the effect of offline temperature change compensation

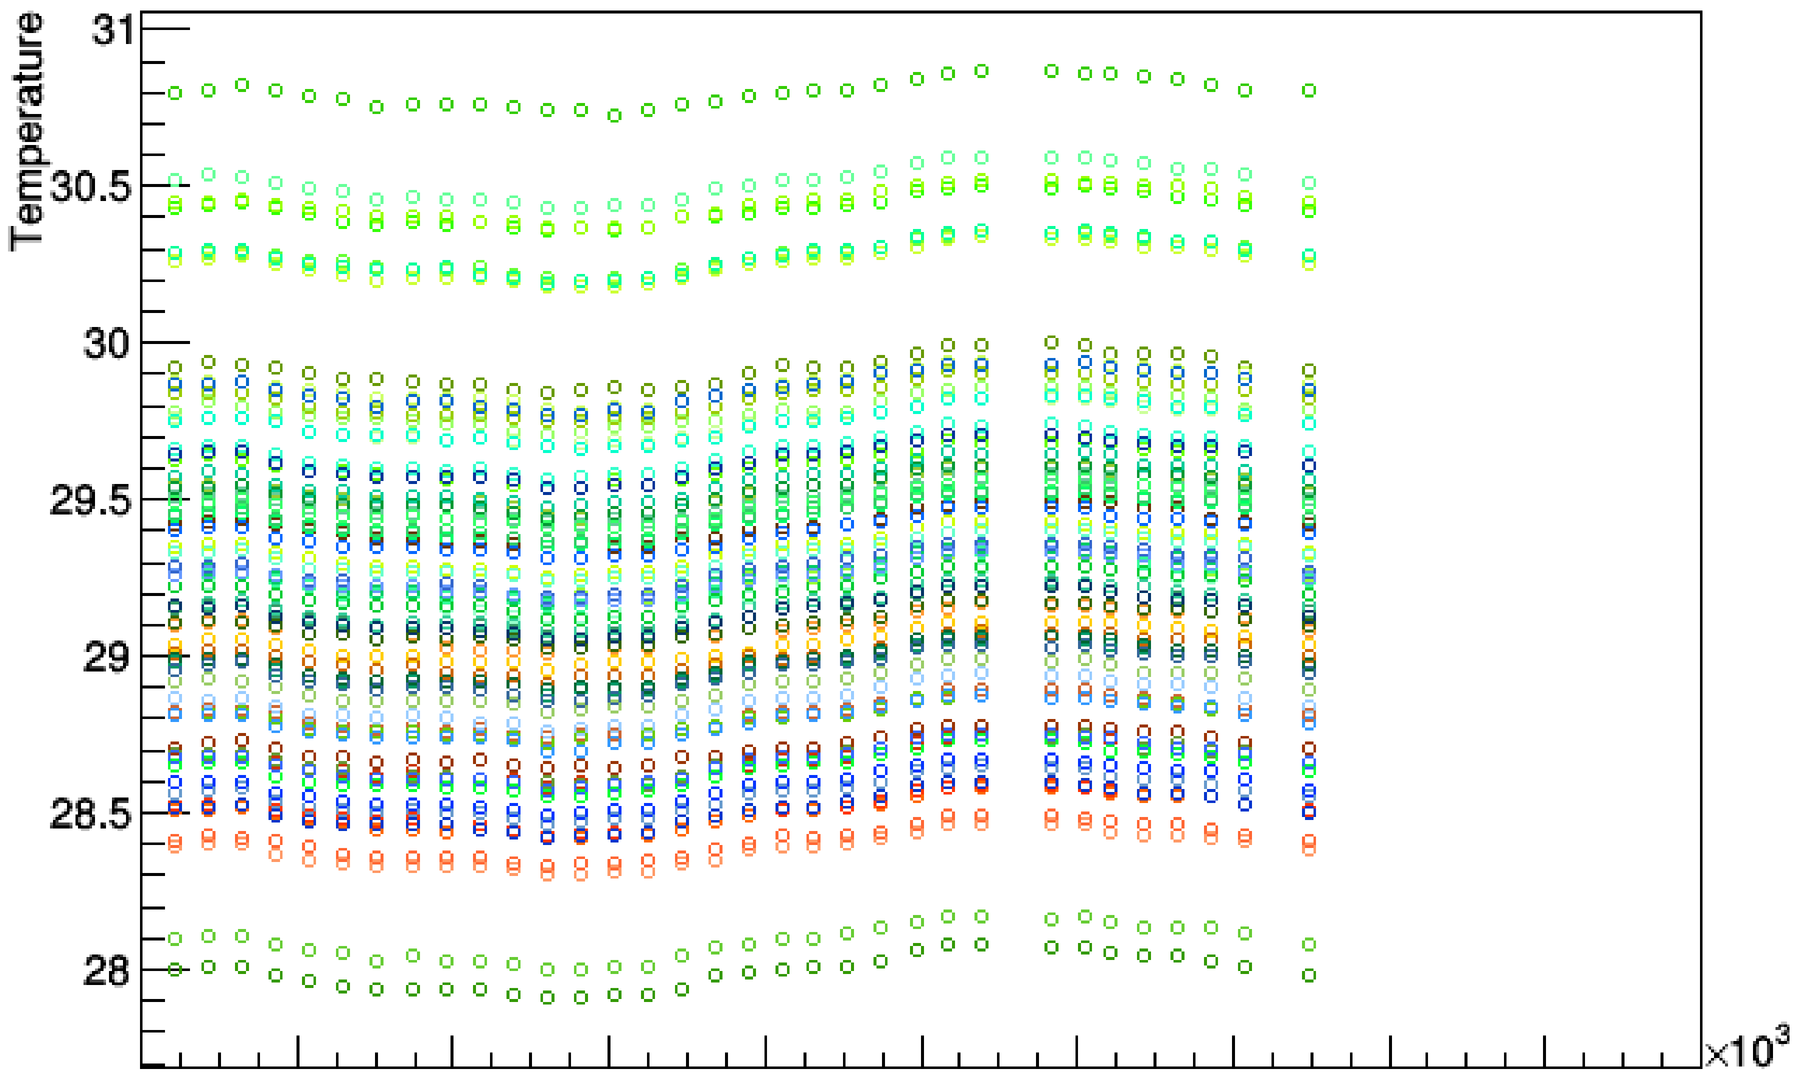

Periodic temperature readings embedded in the raw data

No need to correlate database entries with run/event numbers

Flashing just one test beam result

EMCAL Energy Resolution and Linearity

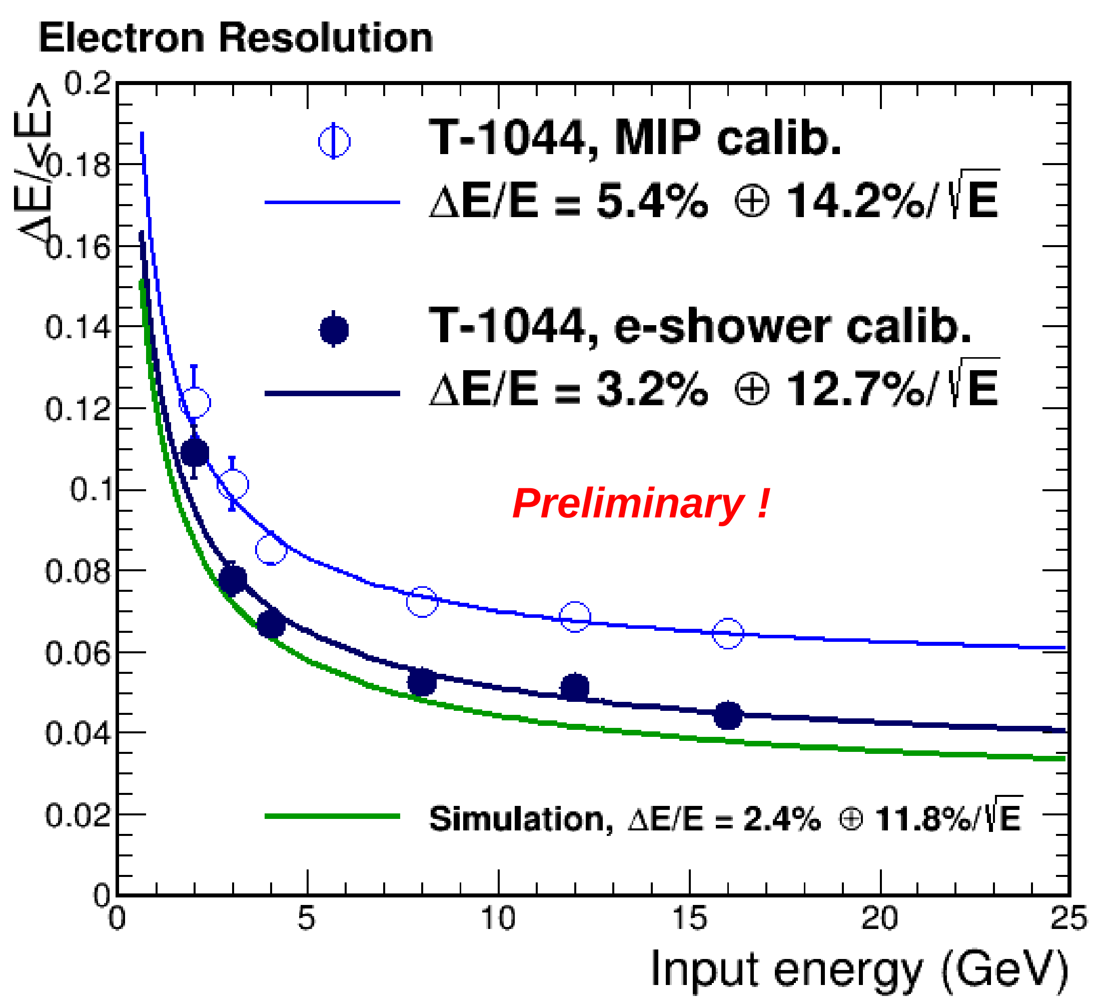

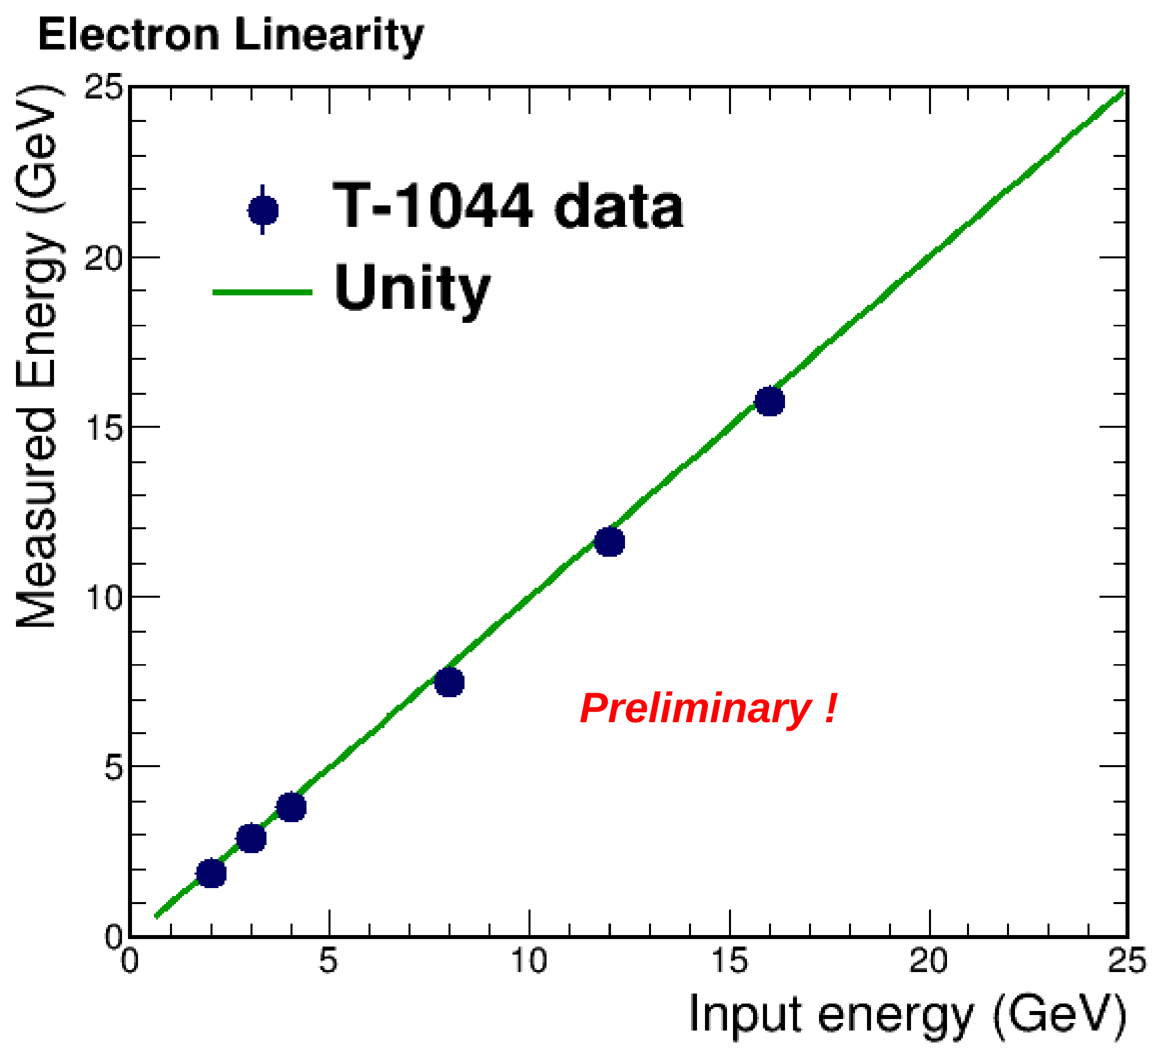

Summary

- sPHENIX will re-use the really modern parts of the PHENIX DAQ
- A workhorse “lean” DAQ system in hand to get us through all R&D
- Uses the actual sPHENIX electronics but can also read out standard devices, CAEN, the SRS, DRS4, etc
- Rich set of metadata logging features (“anything your computer knows”) for lab setups, test beams, R&D in general
- Test bed for new electronics to come
- SiPM automated temperature compensation in progress
- SPHENIX envisioned to see first beam in 2022

Taking a hint from Ikea

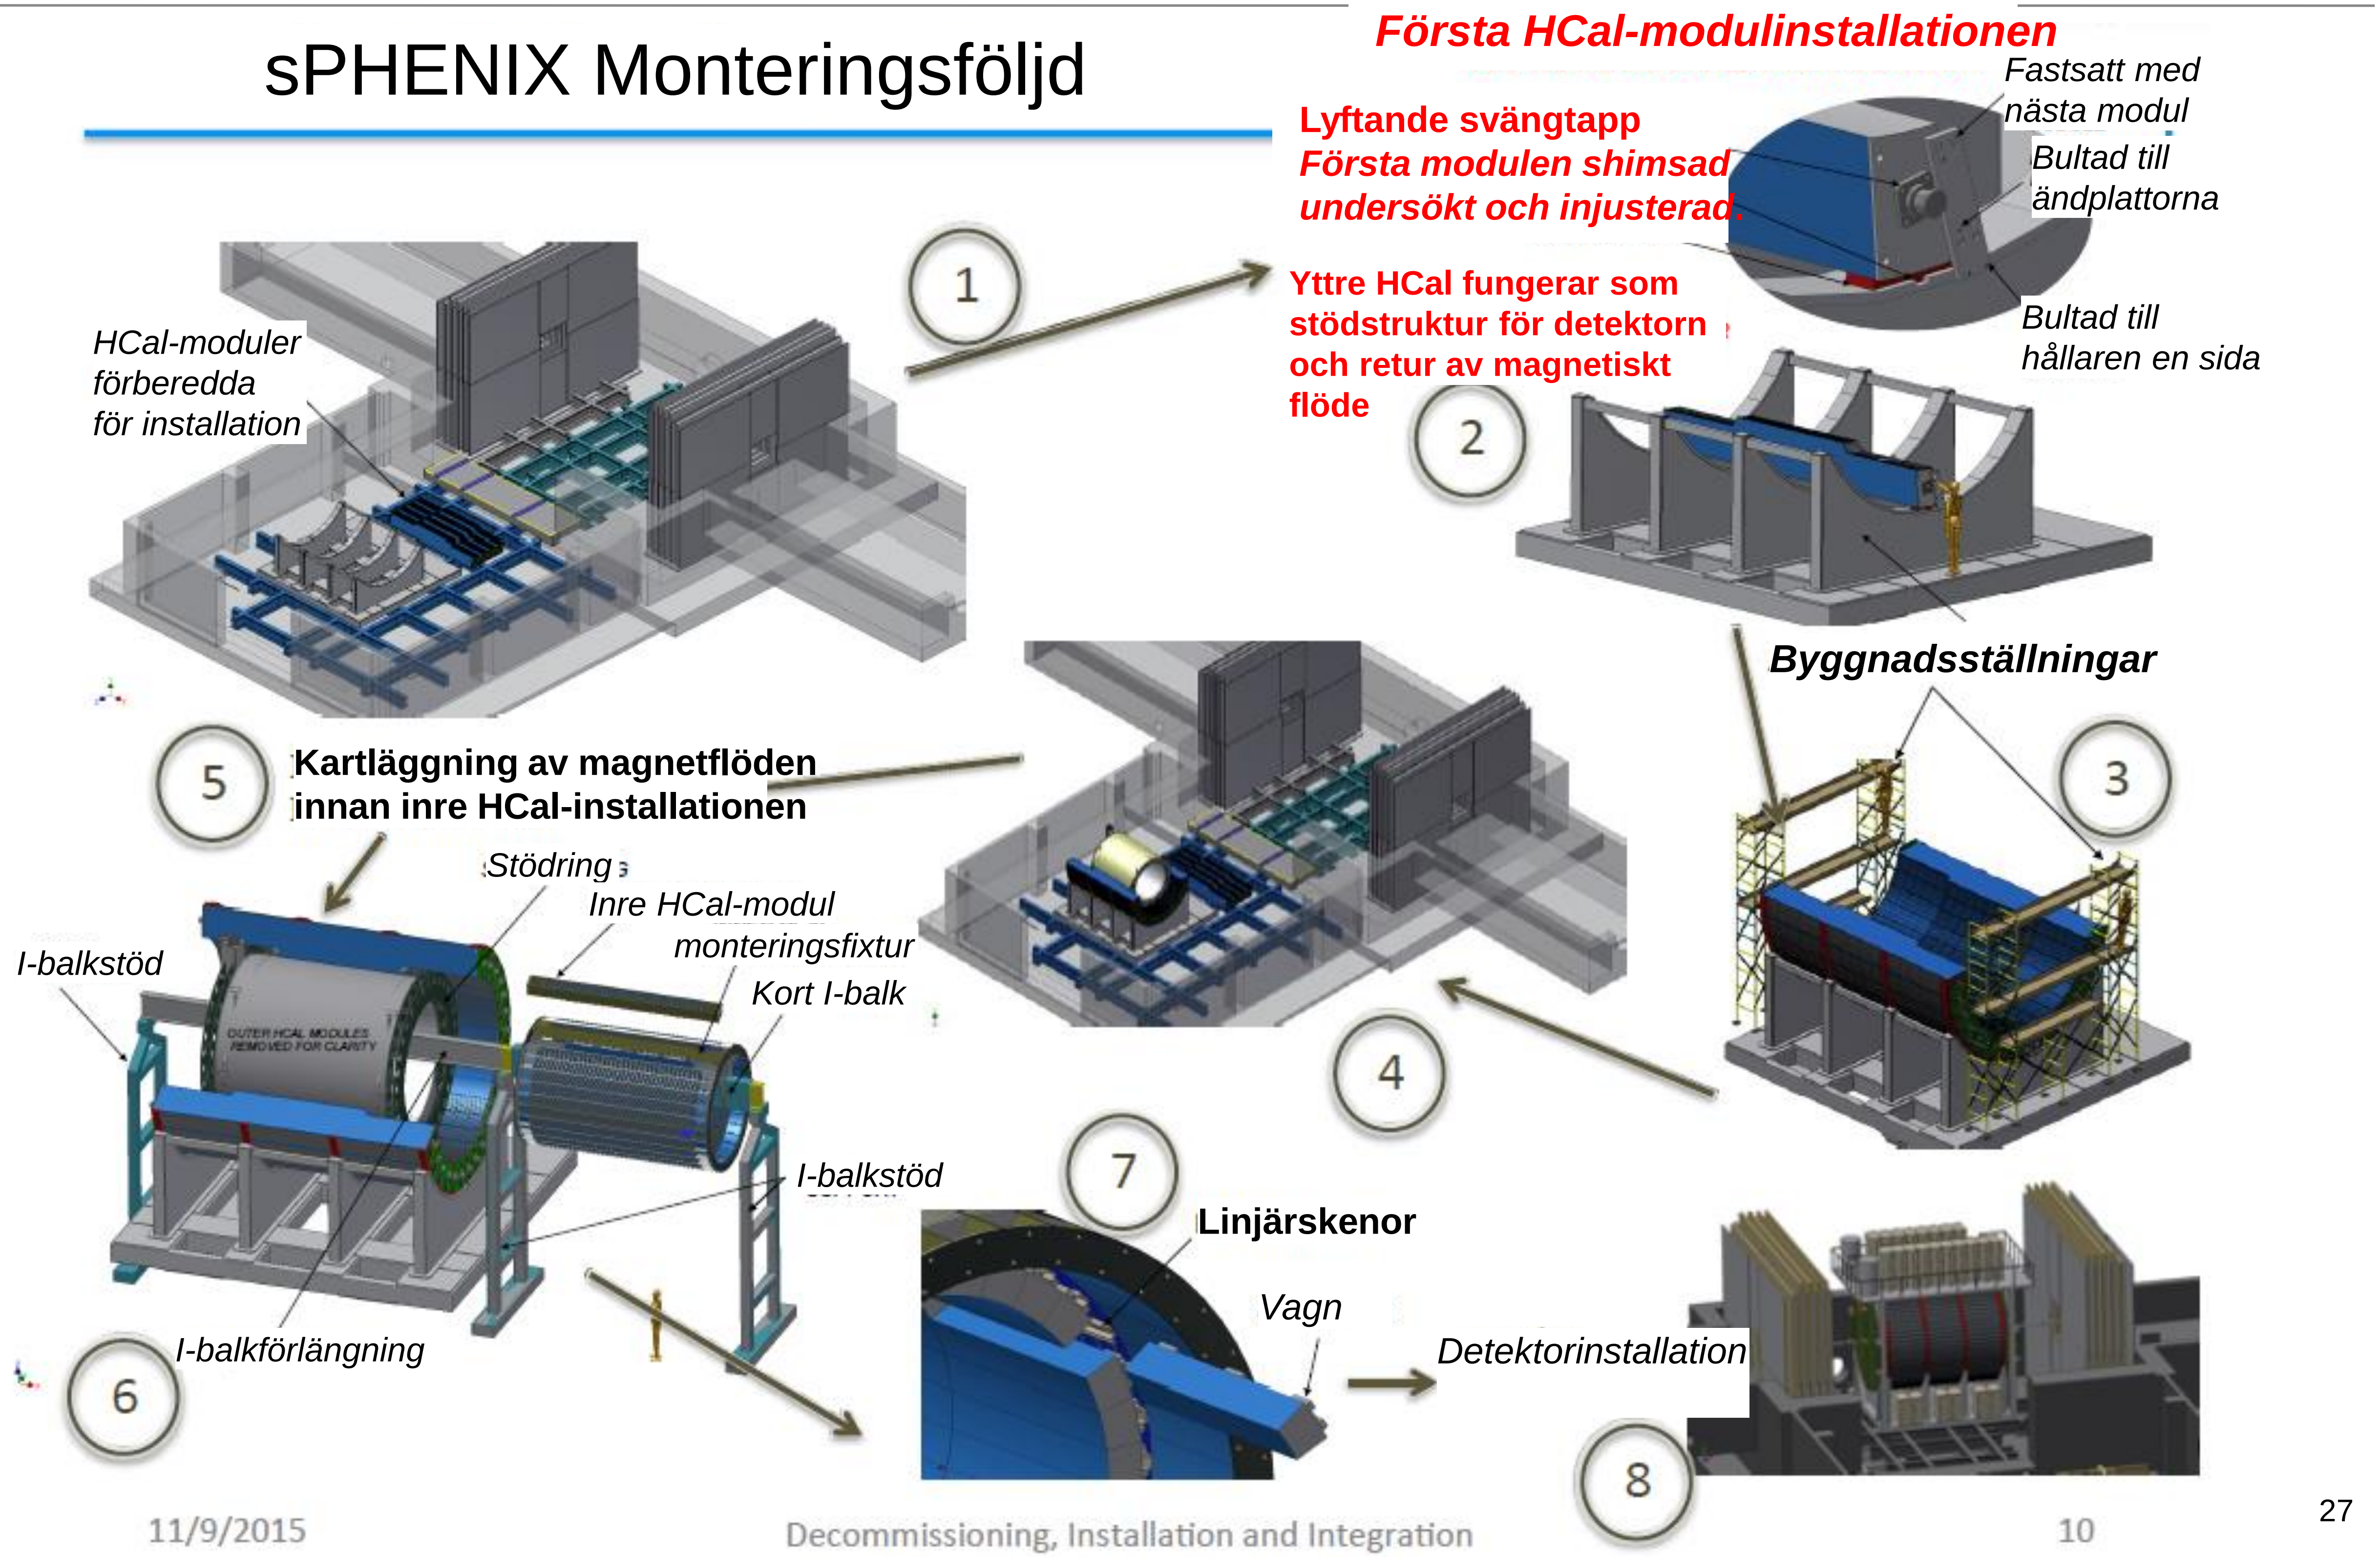

This page is intentionally left blank

DAQ Components – jSEB2 and SEB

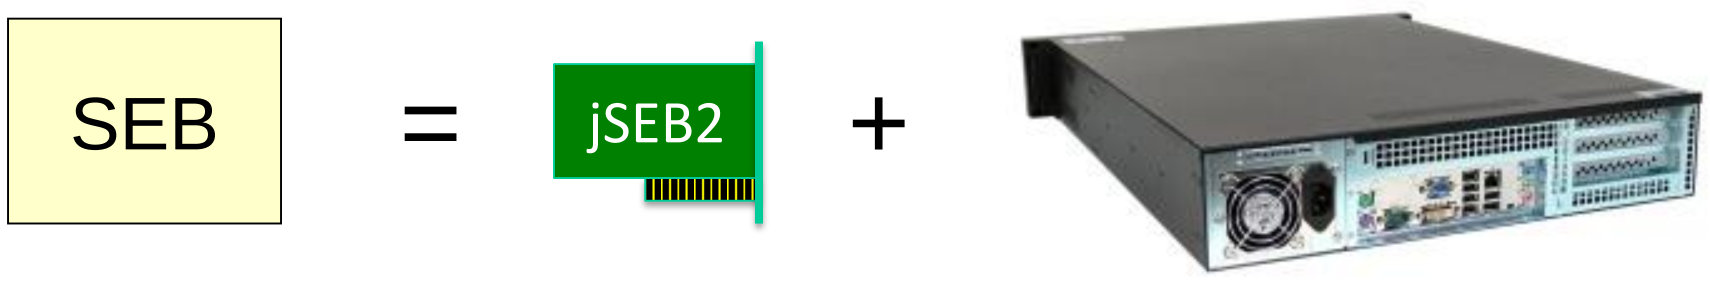

A “Sub-Event Buffer” refers to a Linux PC with a “jSEB2” interface card

Card receives data from the DCM2’s via another board

First time data are seen in a standard PC

jSEB2 is a 4-lane PCI-Express card, 500MByte/s-capable

Code, tools, experience from Runs 14 & 15

Production-grade cards in hand

SEB is basis for test beam, R&D-type, small DAQs

Data compression

After all data *reduction* techniques (zero-suppression, bit-packing, etc) are applied, you typically find that your raw data are still gzip-compressible to a significant amount

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

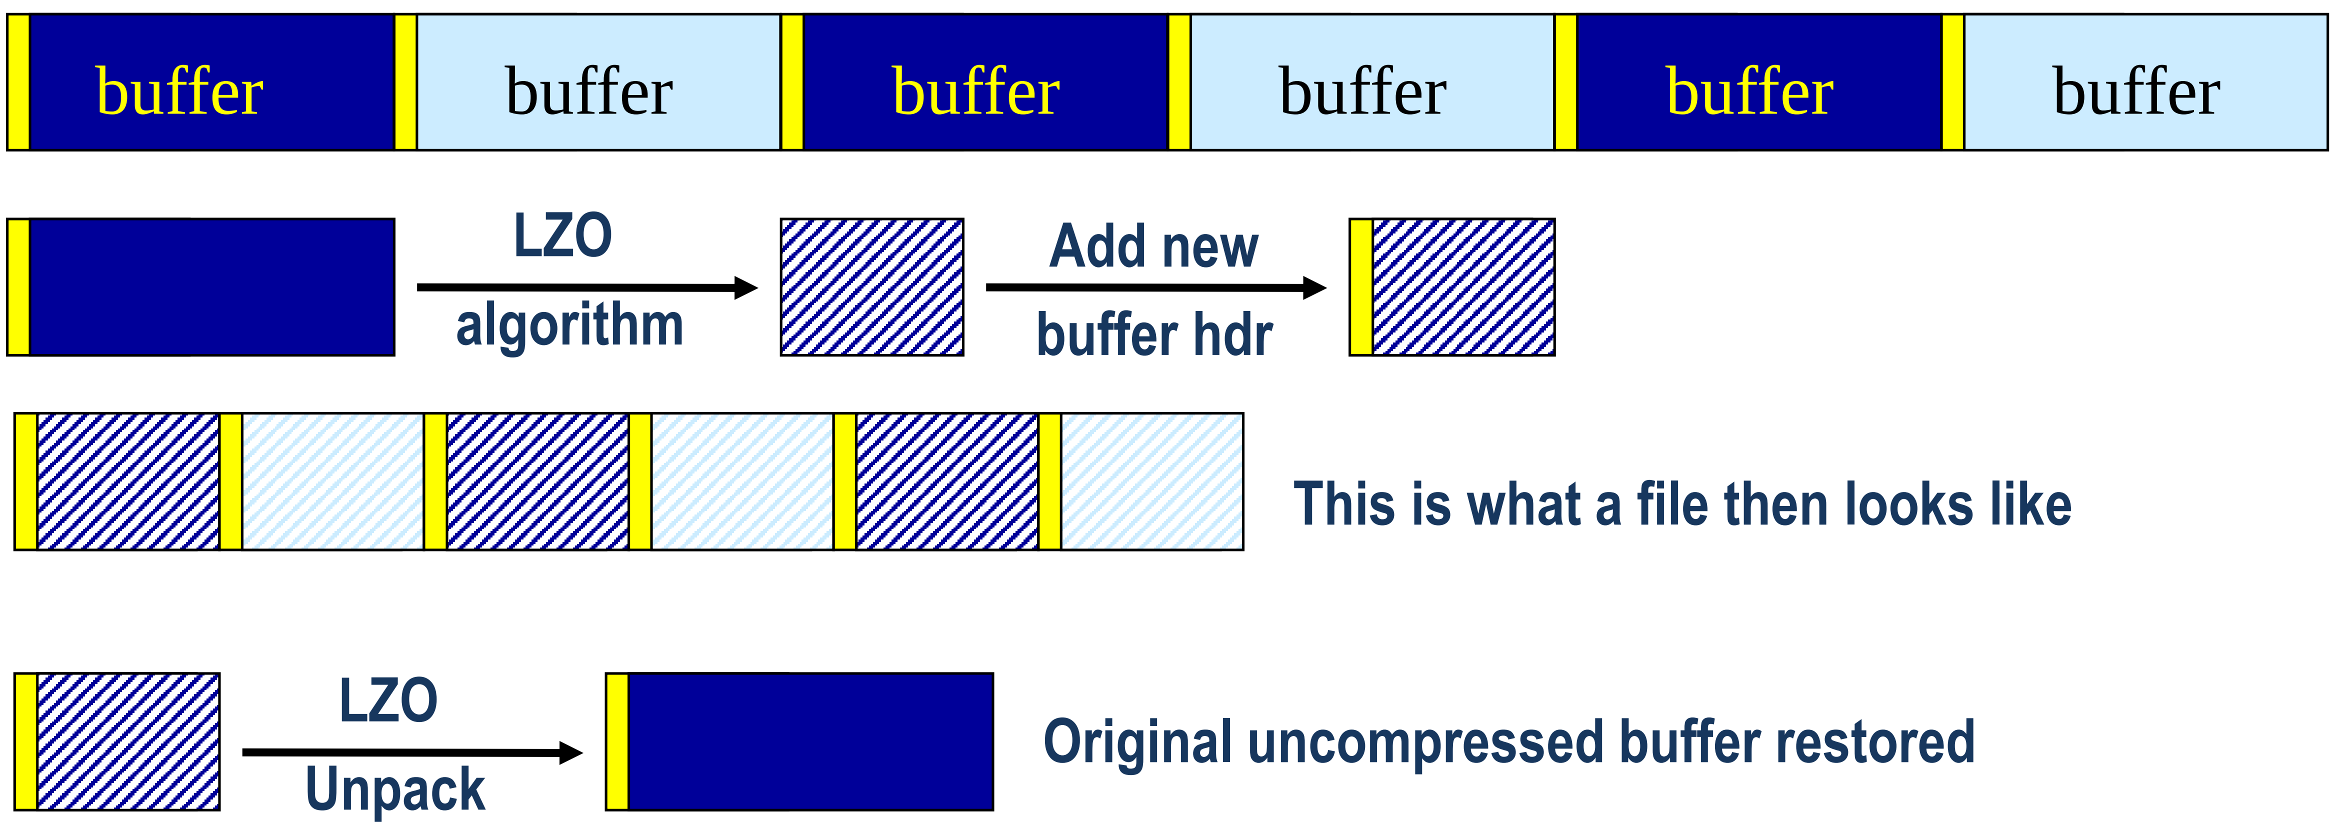

All this is handled completely in the I/O layer, the higher-level routines just receive a buffer as before.

Data compression load

This compression buys you a lot spare bandwidth and nice features (e.g, your analysis runs faster!)

However:

A late-stage compression (think: in your tape drive in the extreme case) doesn't really help you

The compression has to kick in before the data hit your storage system for the first time

No single machine can keep up with compressing a >2GByte/s rate

Switch to distributed compression

The Event builder has to cope with the uncompressed data flow, e.g. 1200MB/s … 2500MB/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

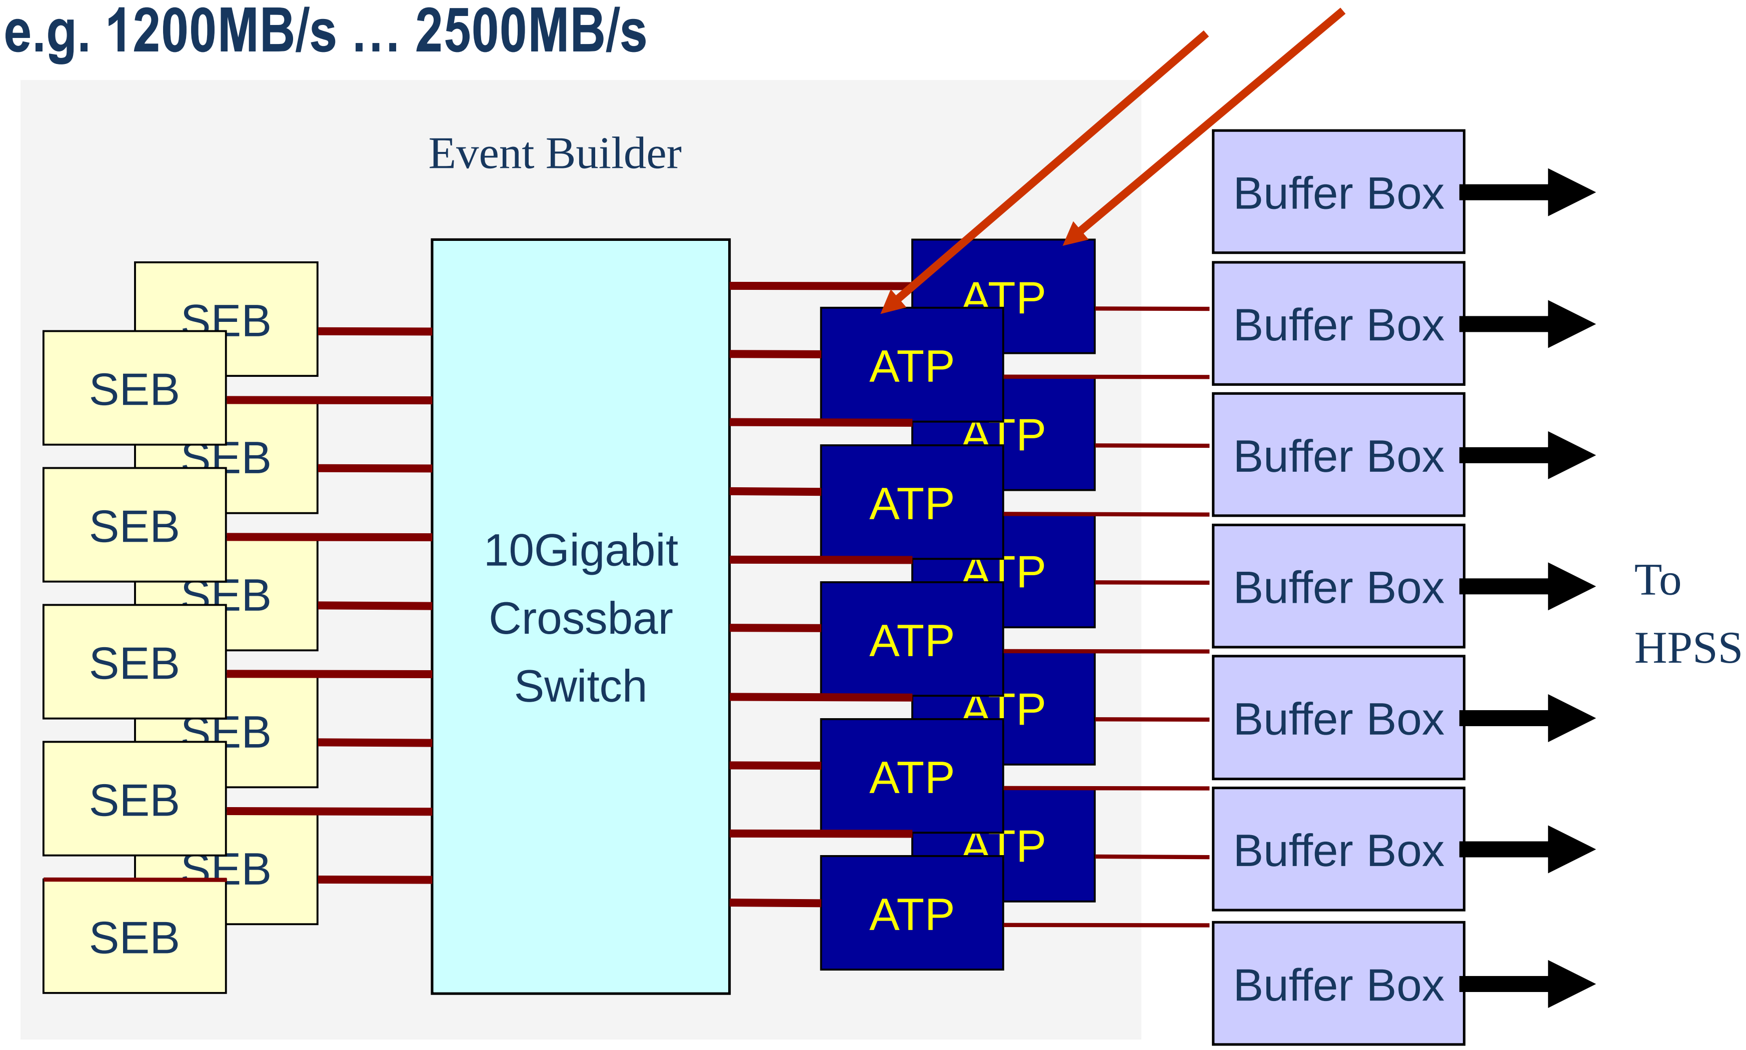

The buffer boxes and storage system see the compressed data stream, 700MB/s … 1300MB/s

Current compression levels:
100G become ~55G

Data Rates (Run14)

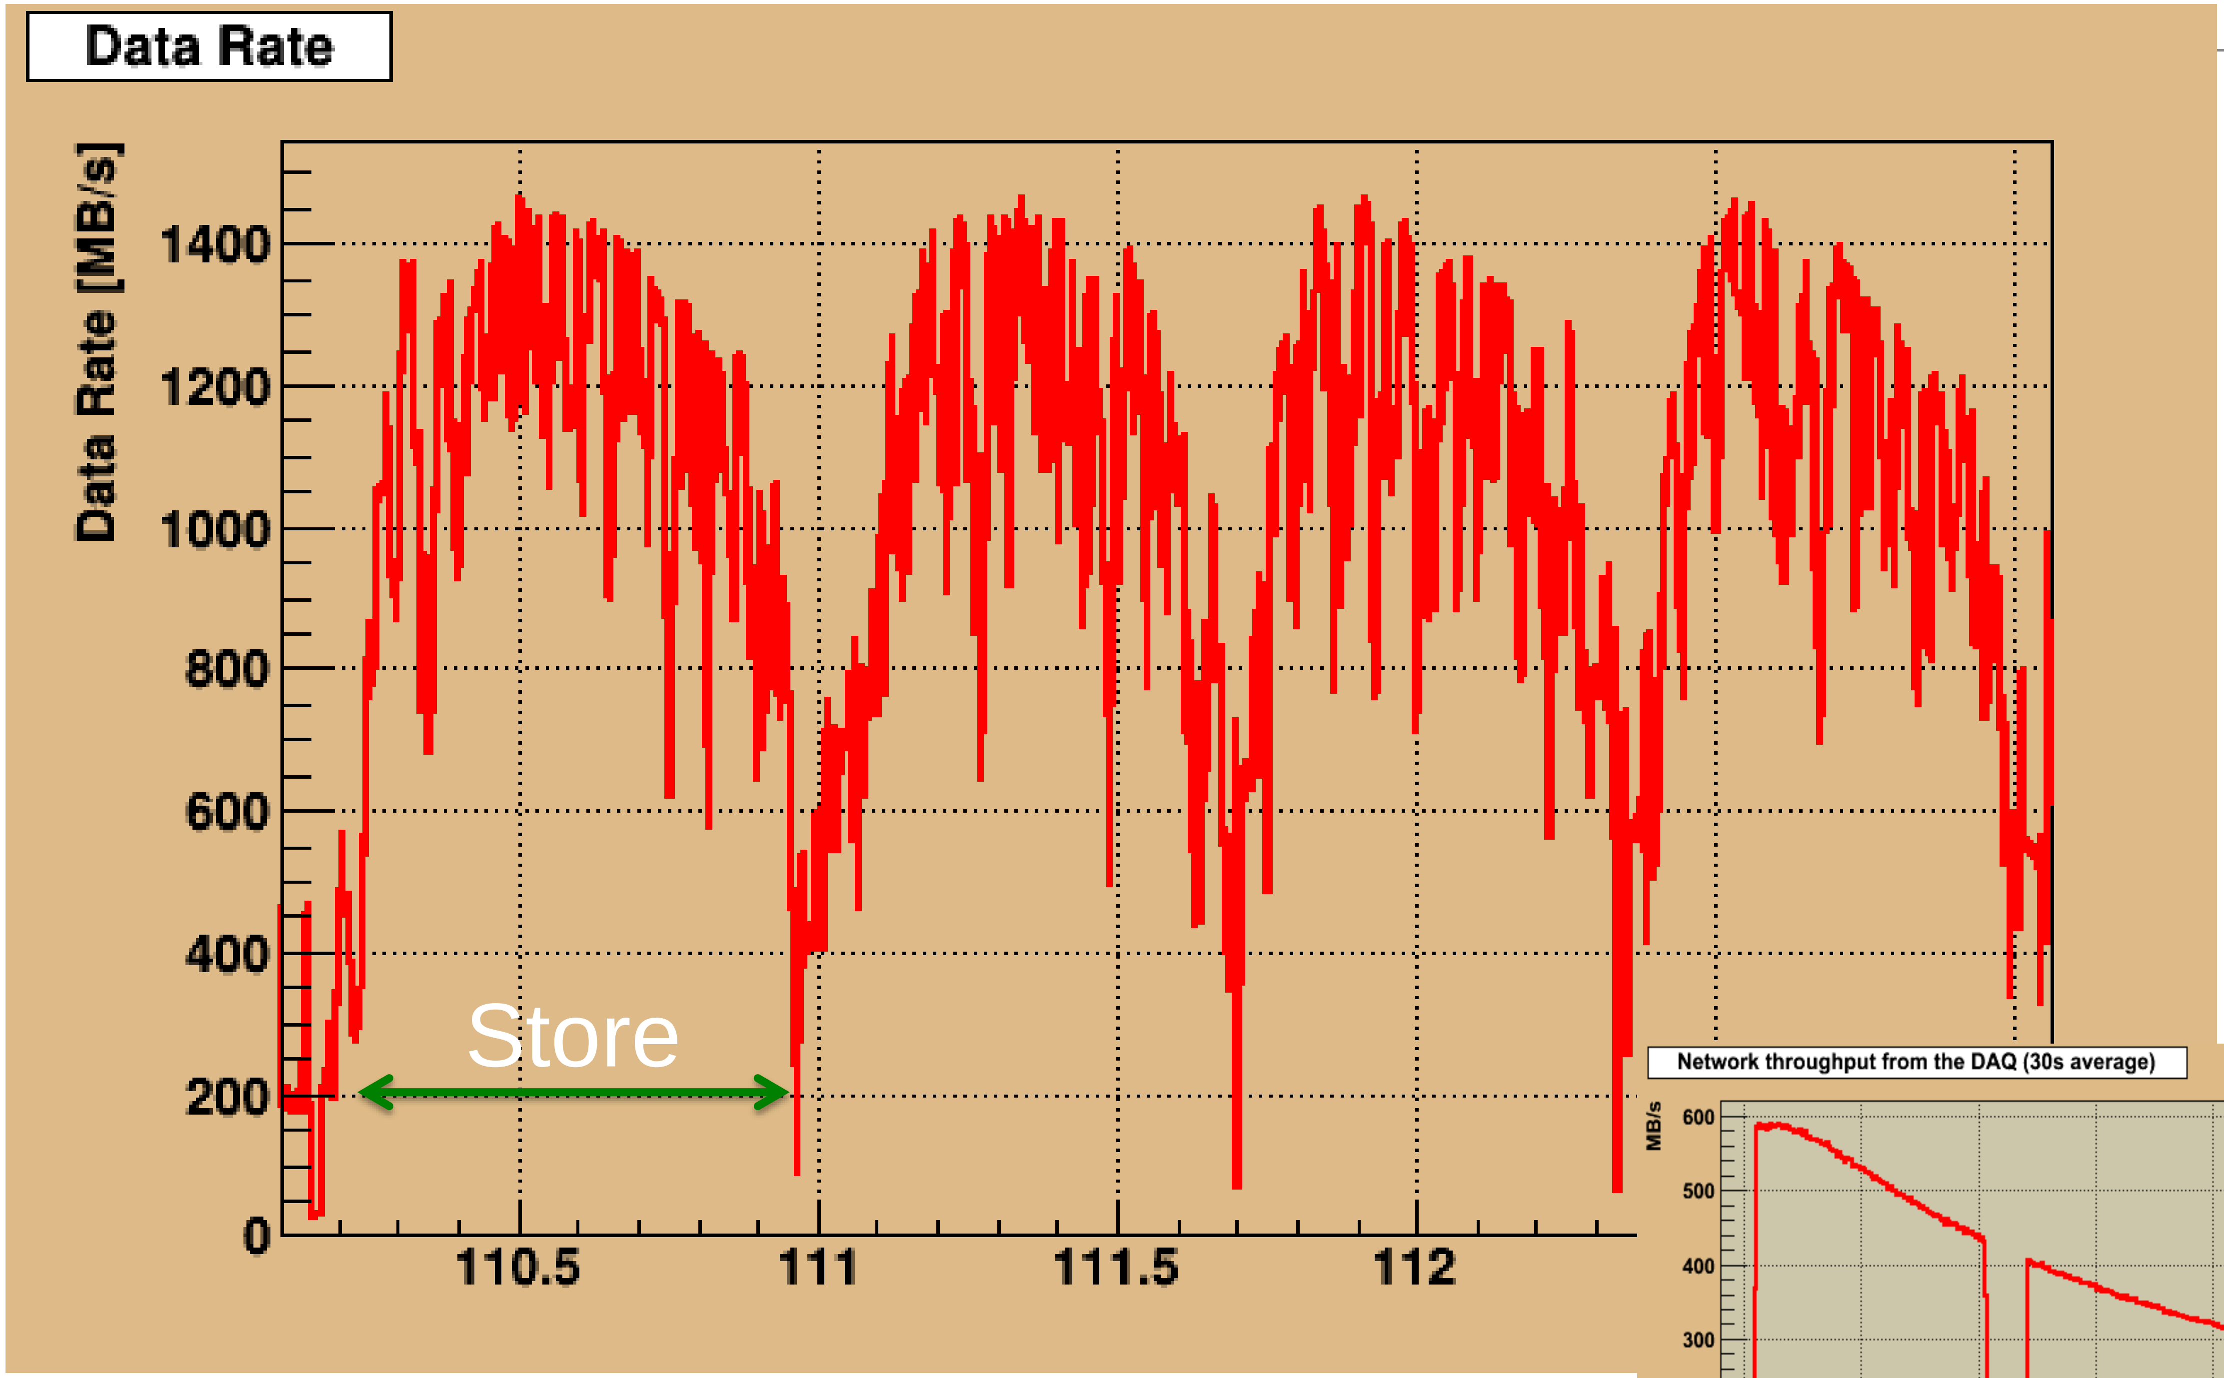

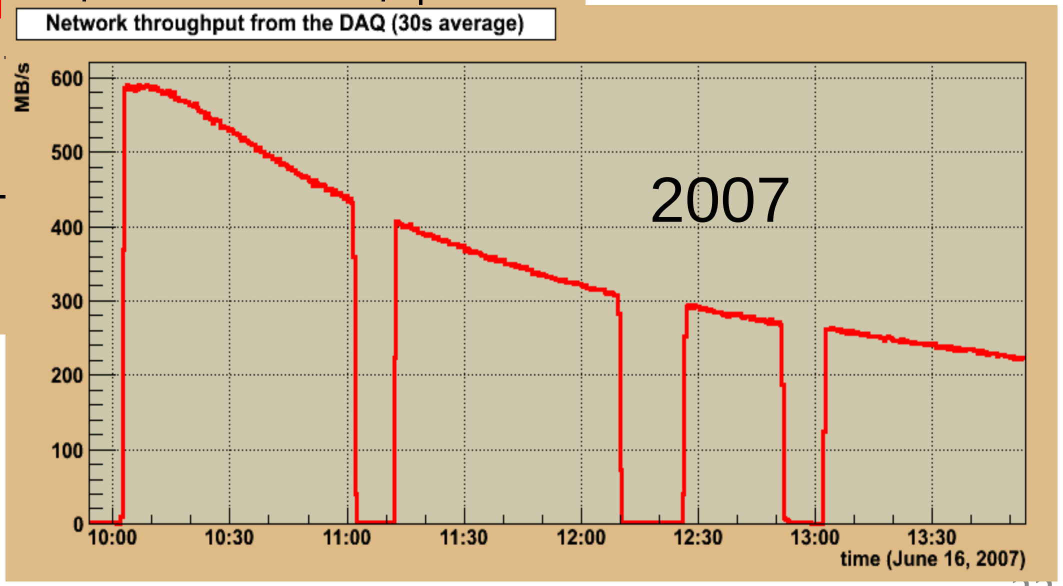

Trigger Electronics & Timing system

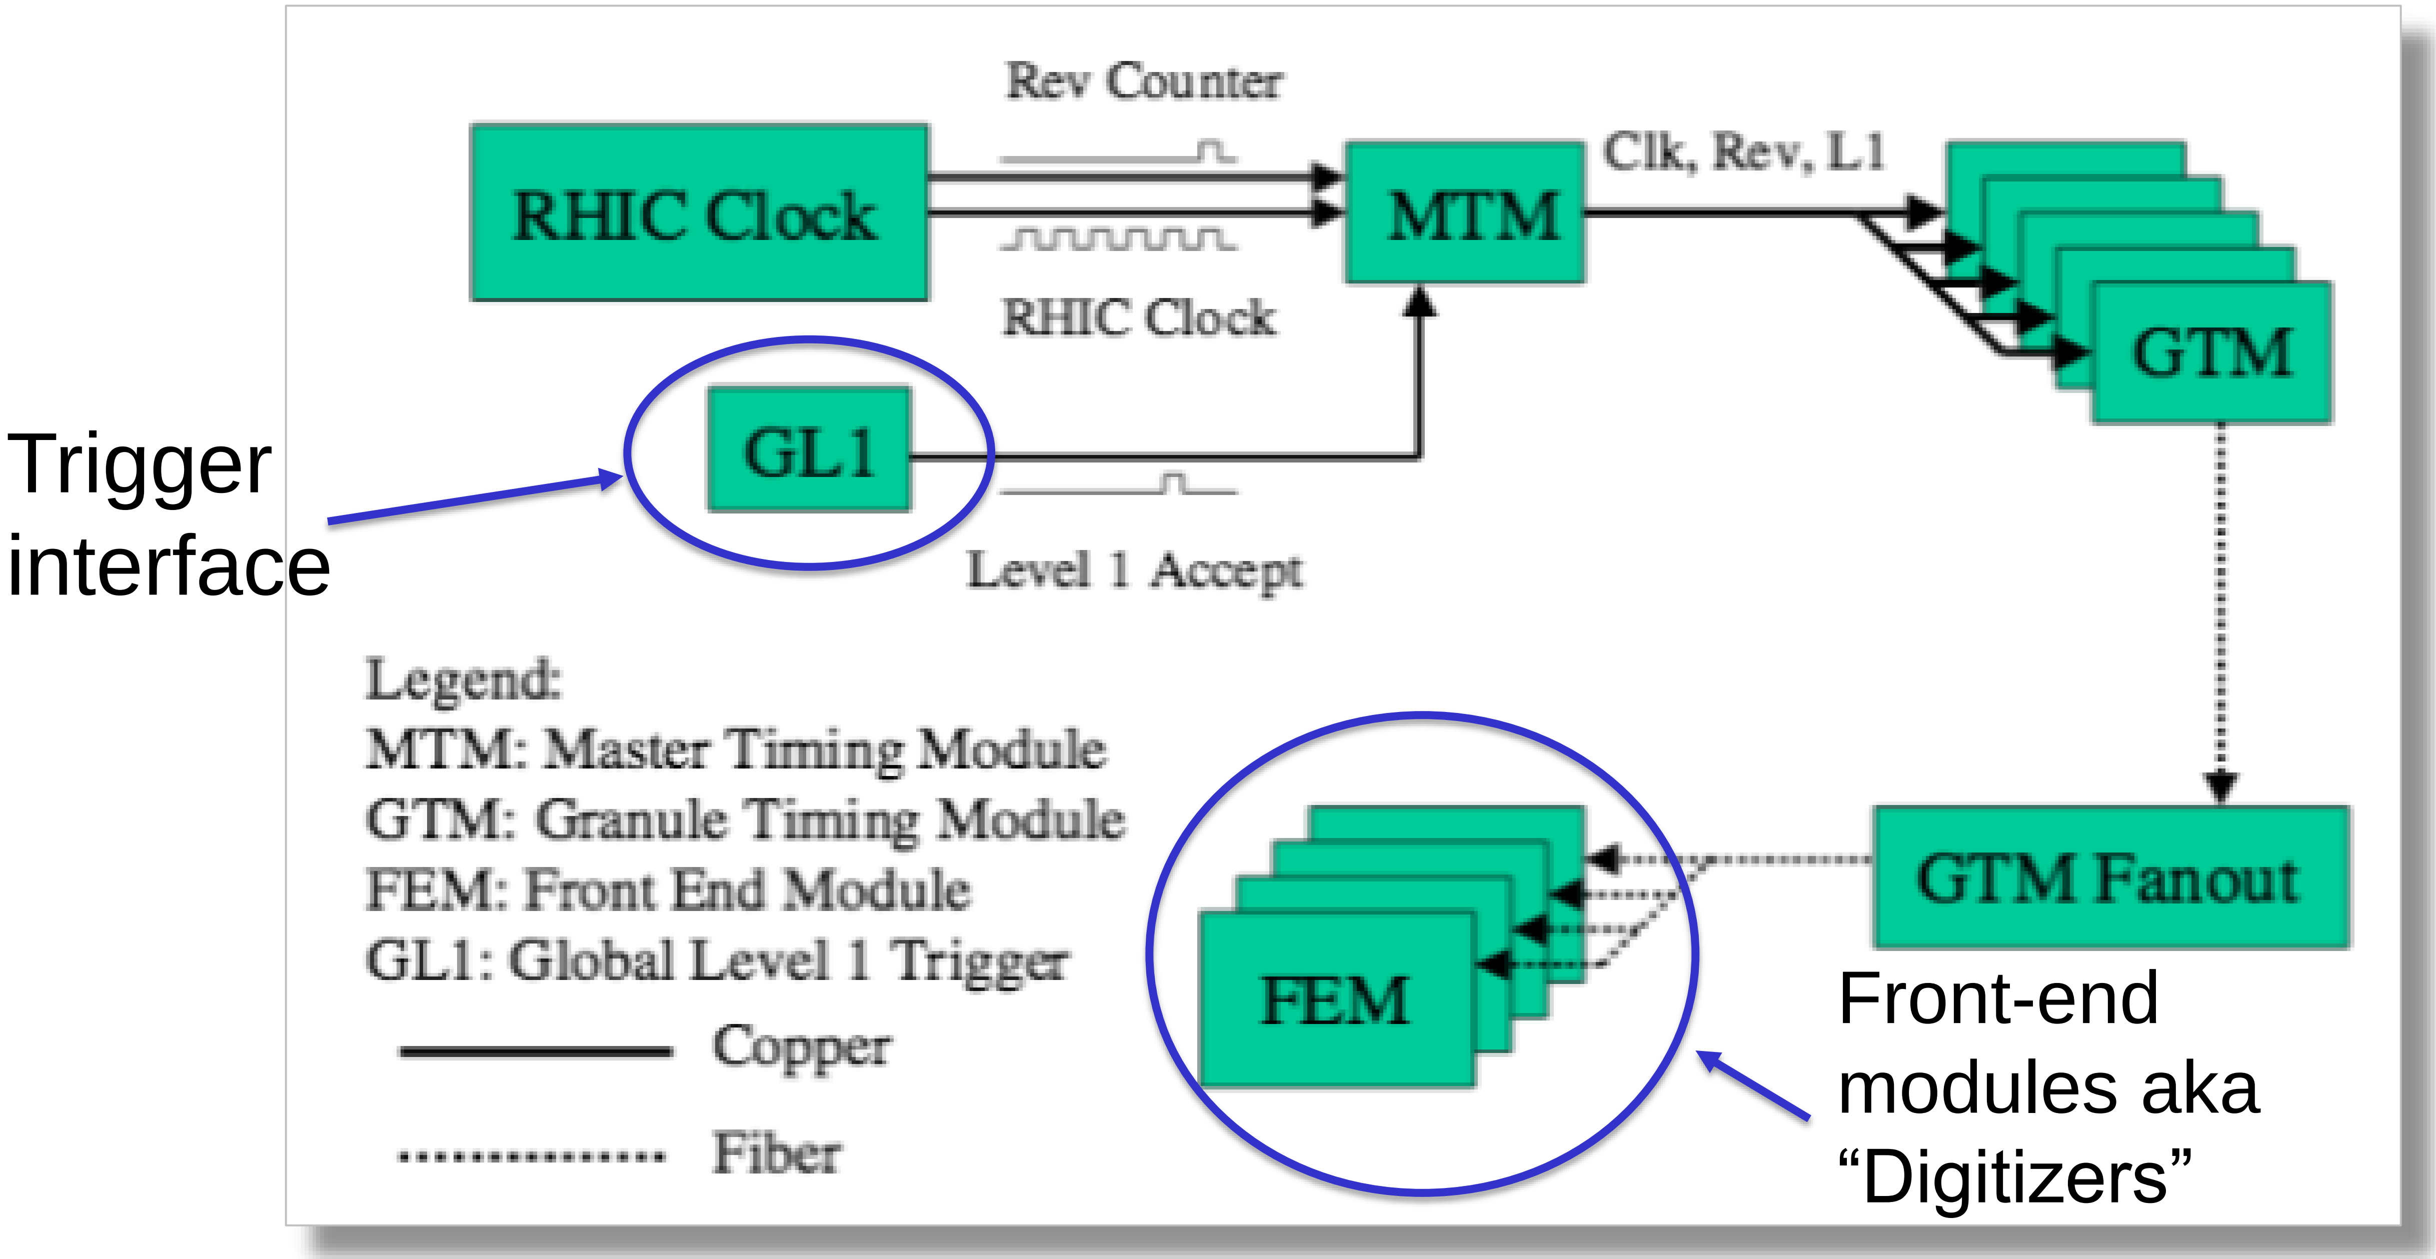

Local-Level 1 (LL1) Boards

- LL1's generate a trigger signal from trigger primitives generated in the front-end for a given detector
- new LL1 boards for calorimeter triggers
- Potential new triggers for p+p, p+A need LL1 boards.

Trigger & Timing system

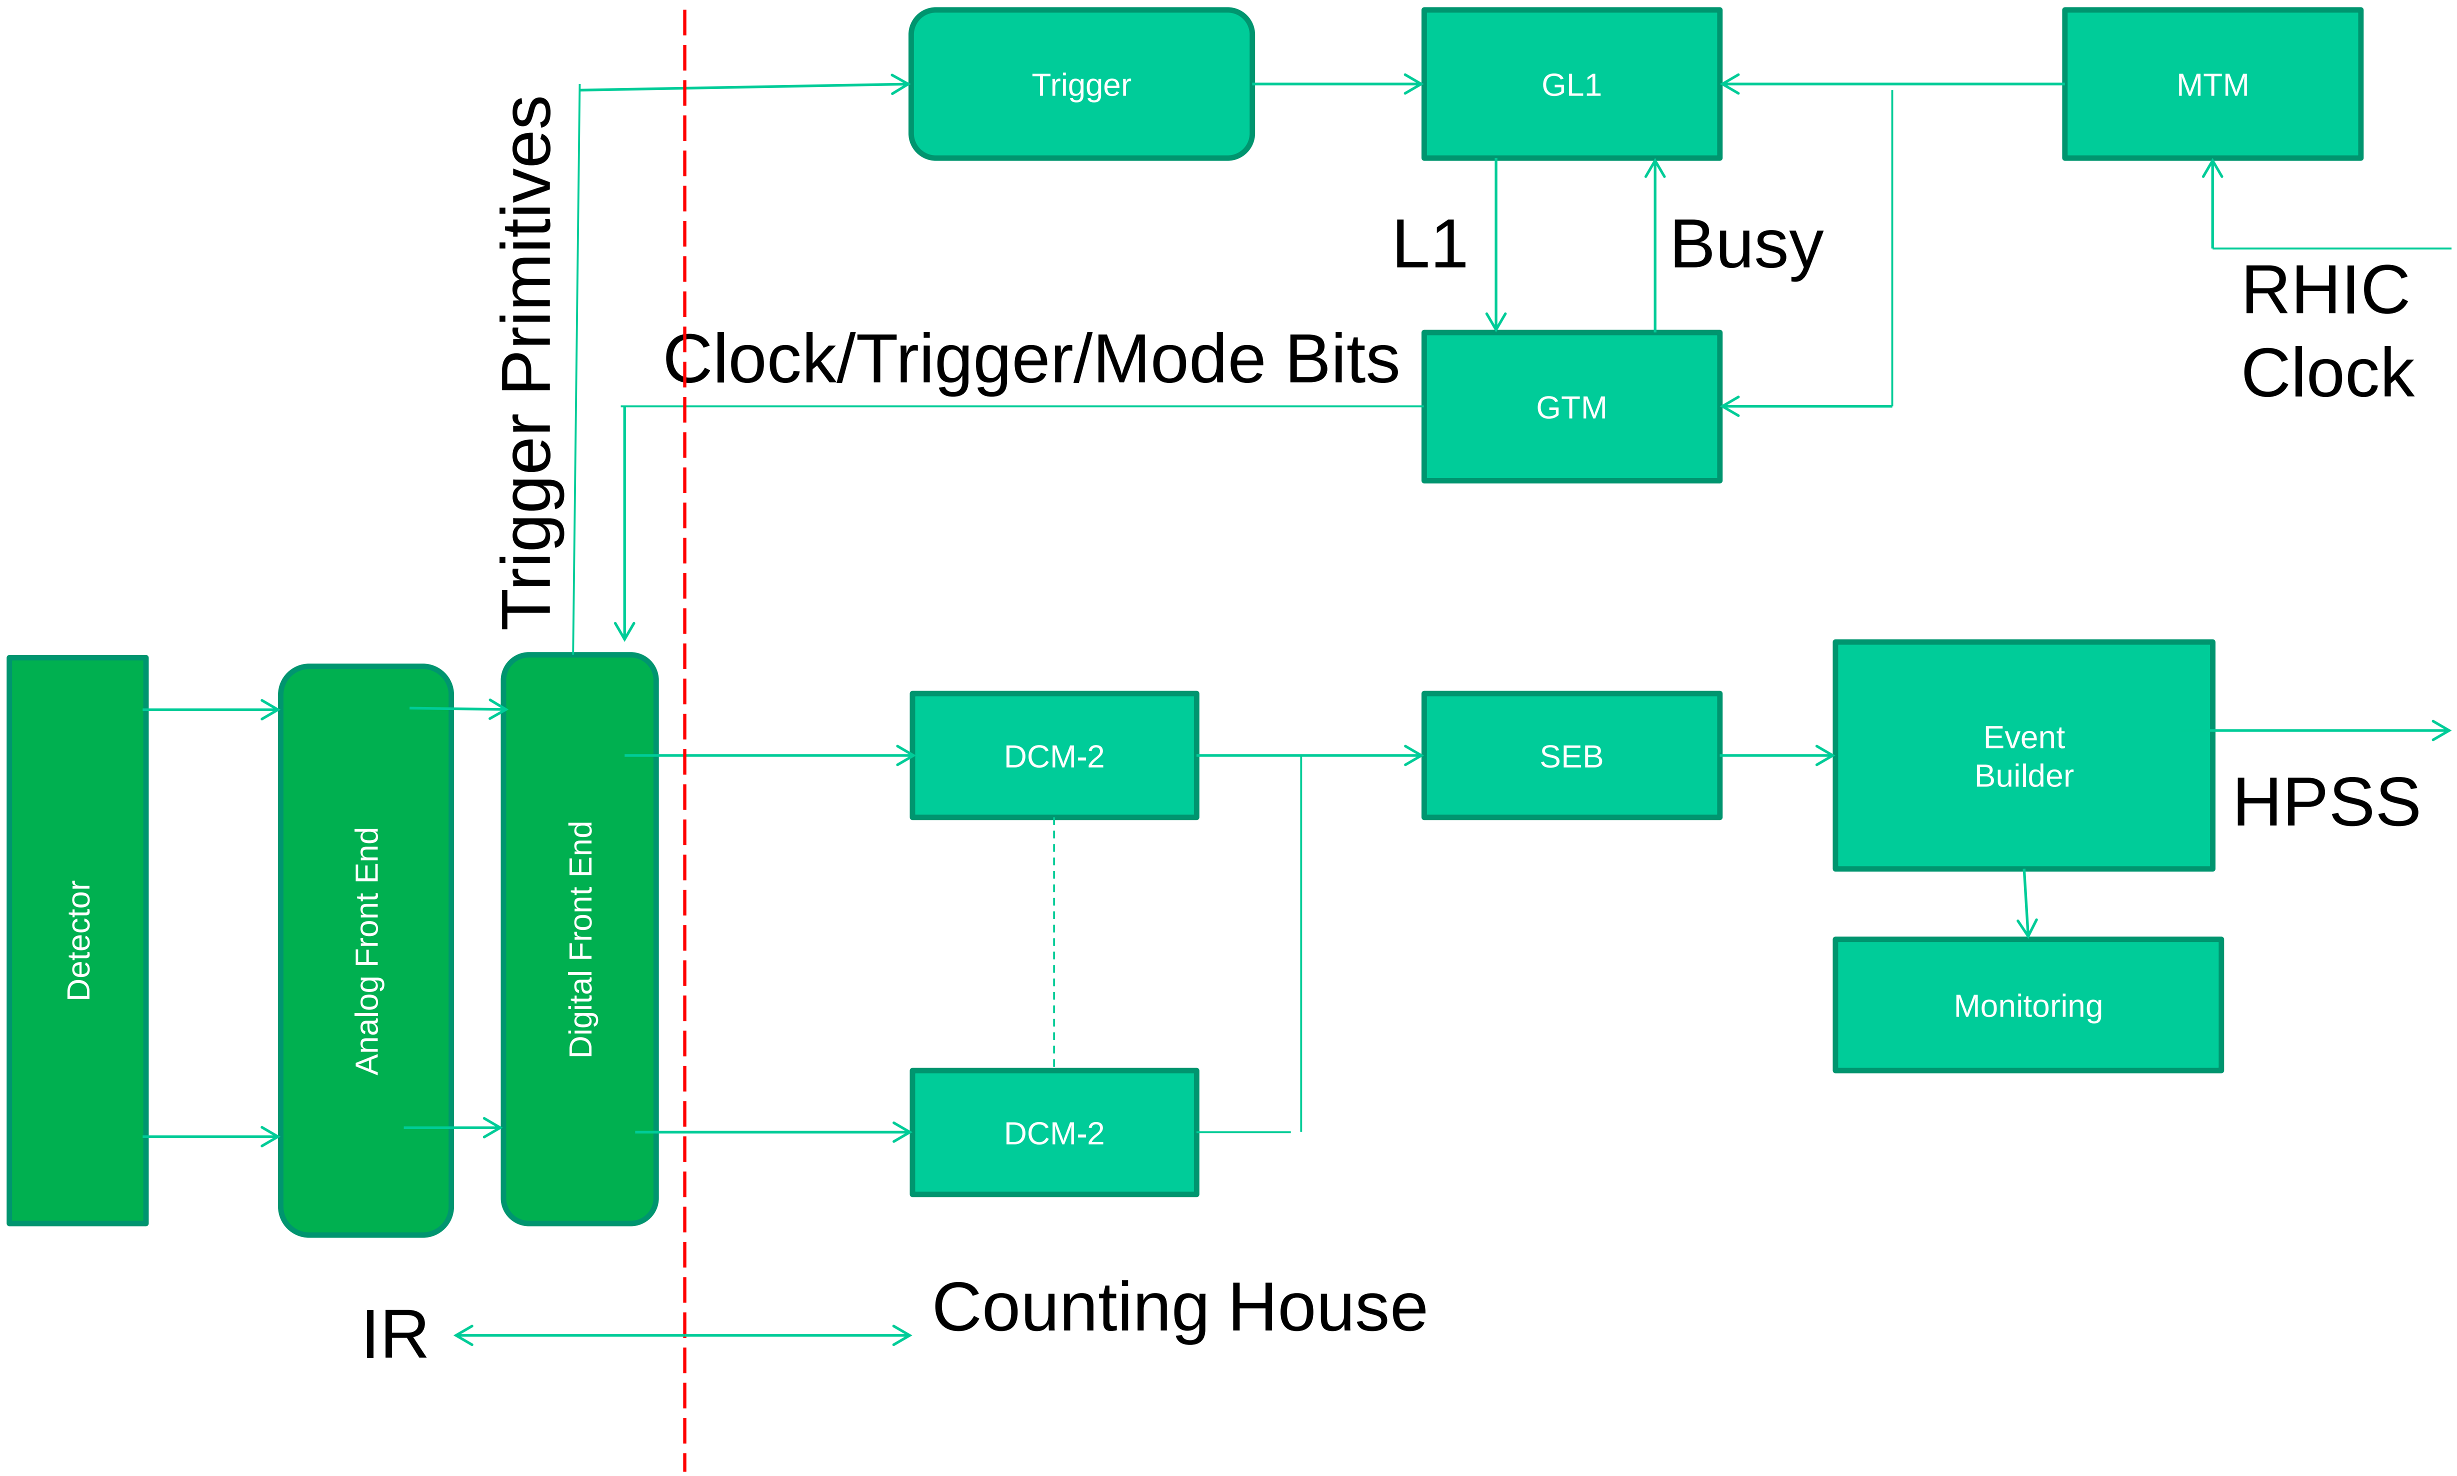

Some envisioned performance plots

The yield of hard processes in a typical RHIC year (22 weeks of running)

The lowest count in the plot correponds to 10 billion events

(We expect to get 20)

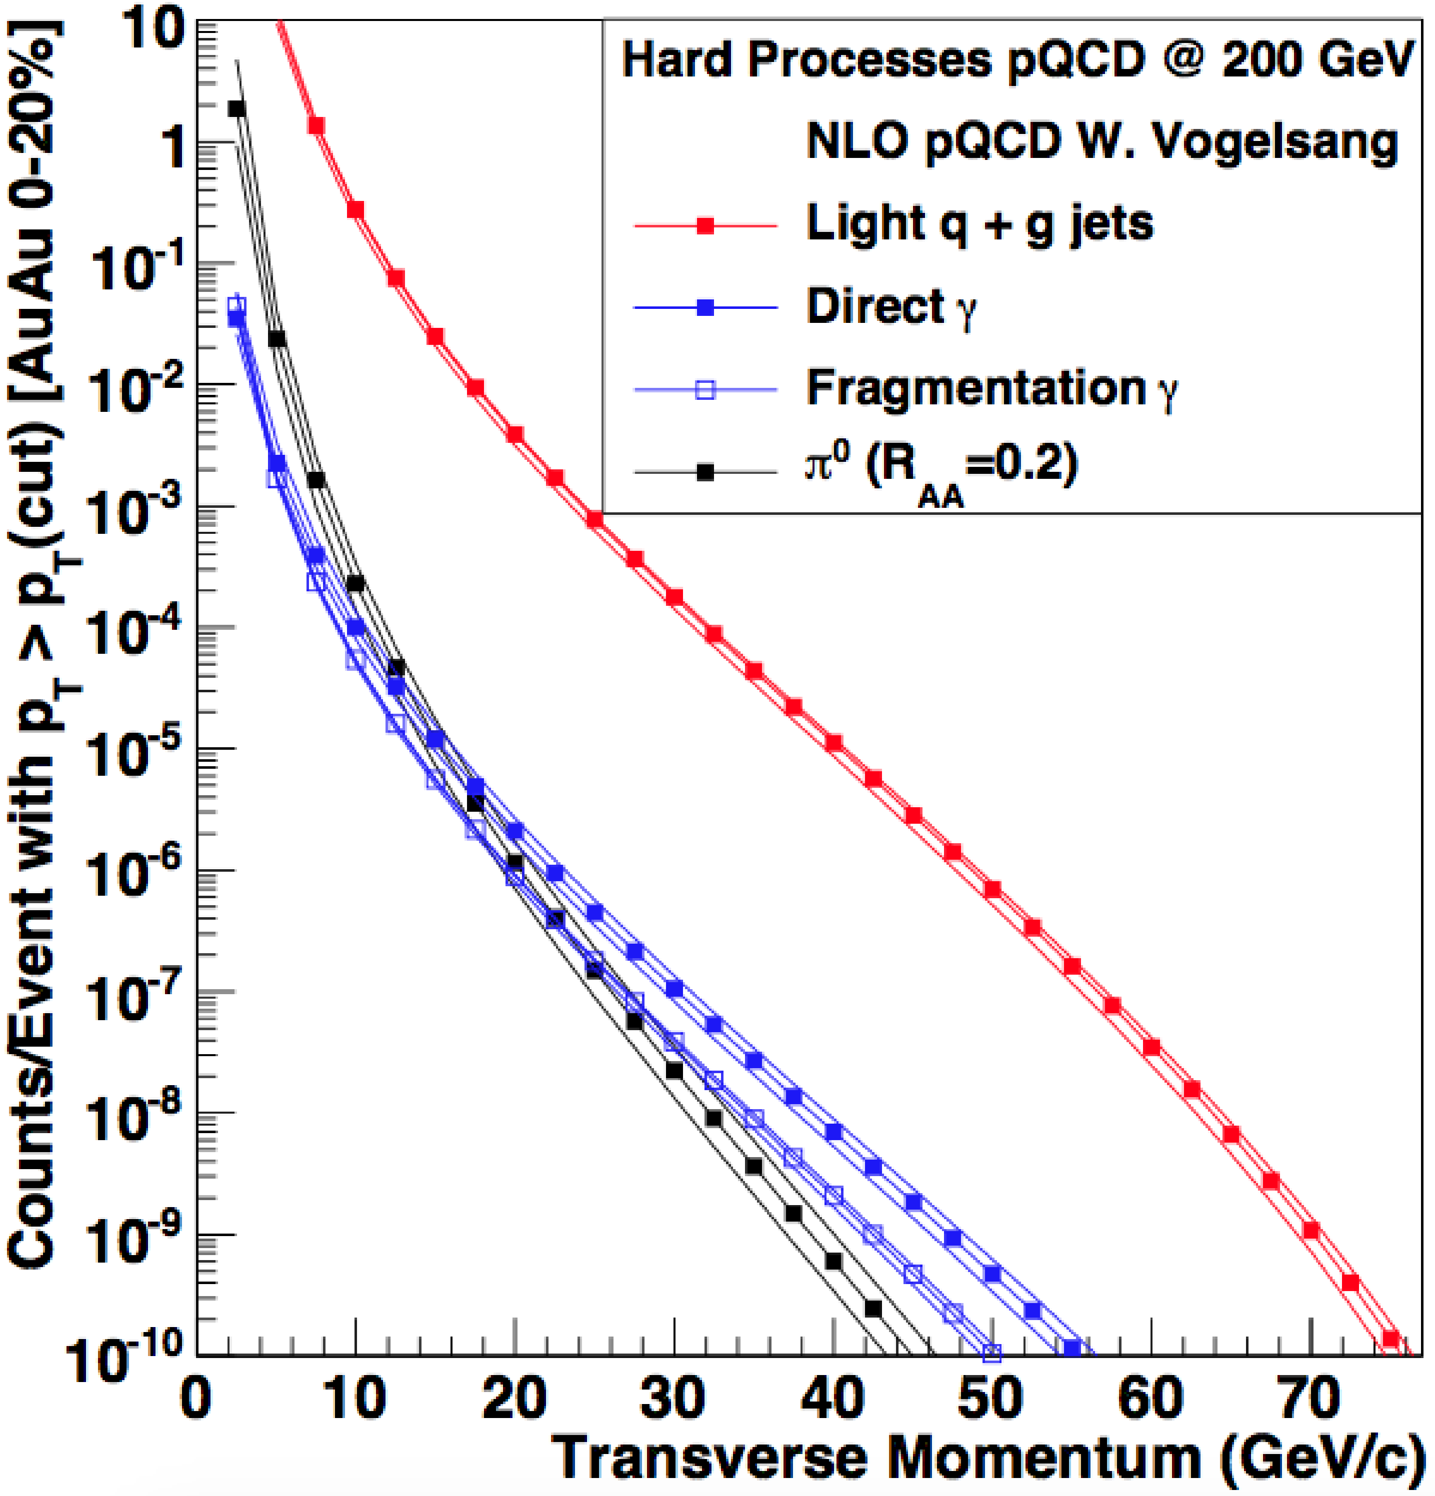

Statistical reach for some probes

The way PHENIX gets to high PT is through π0s

That reaches to ~20GeV/c

In sPHENIX, we can get to 4x that with jets

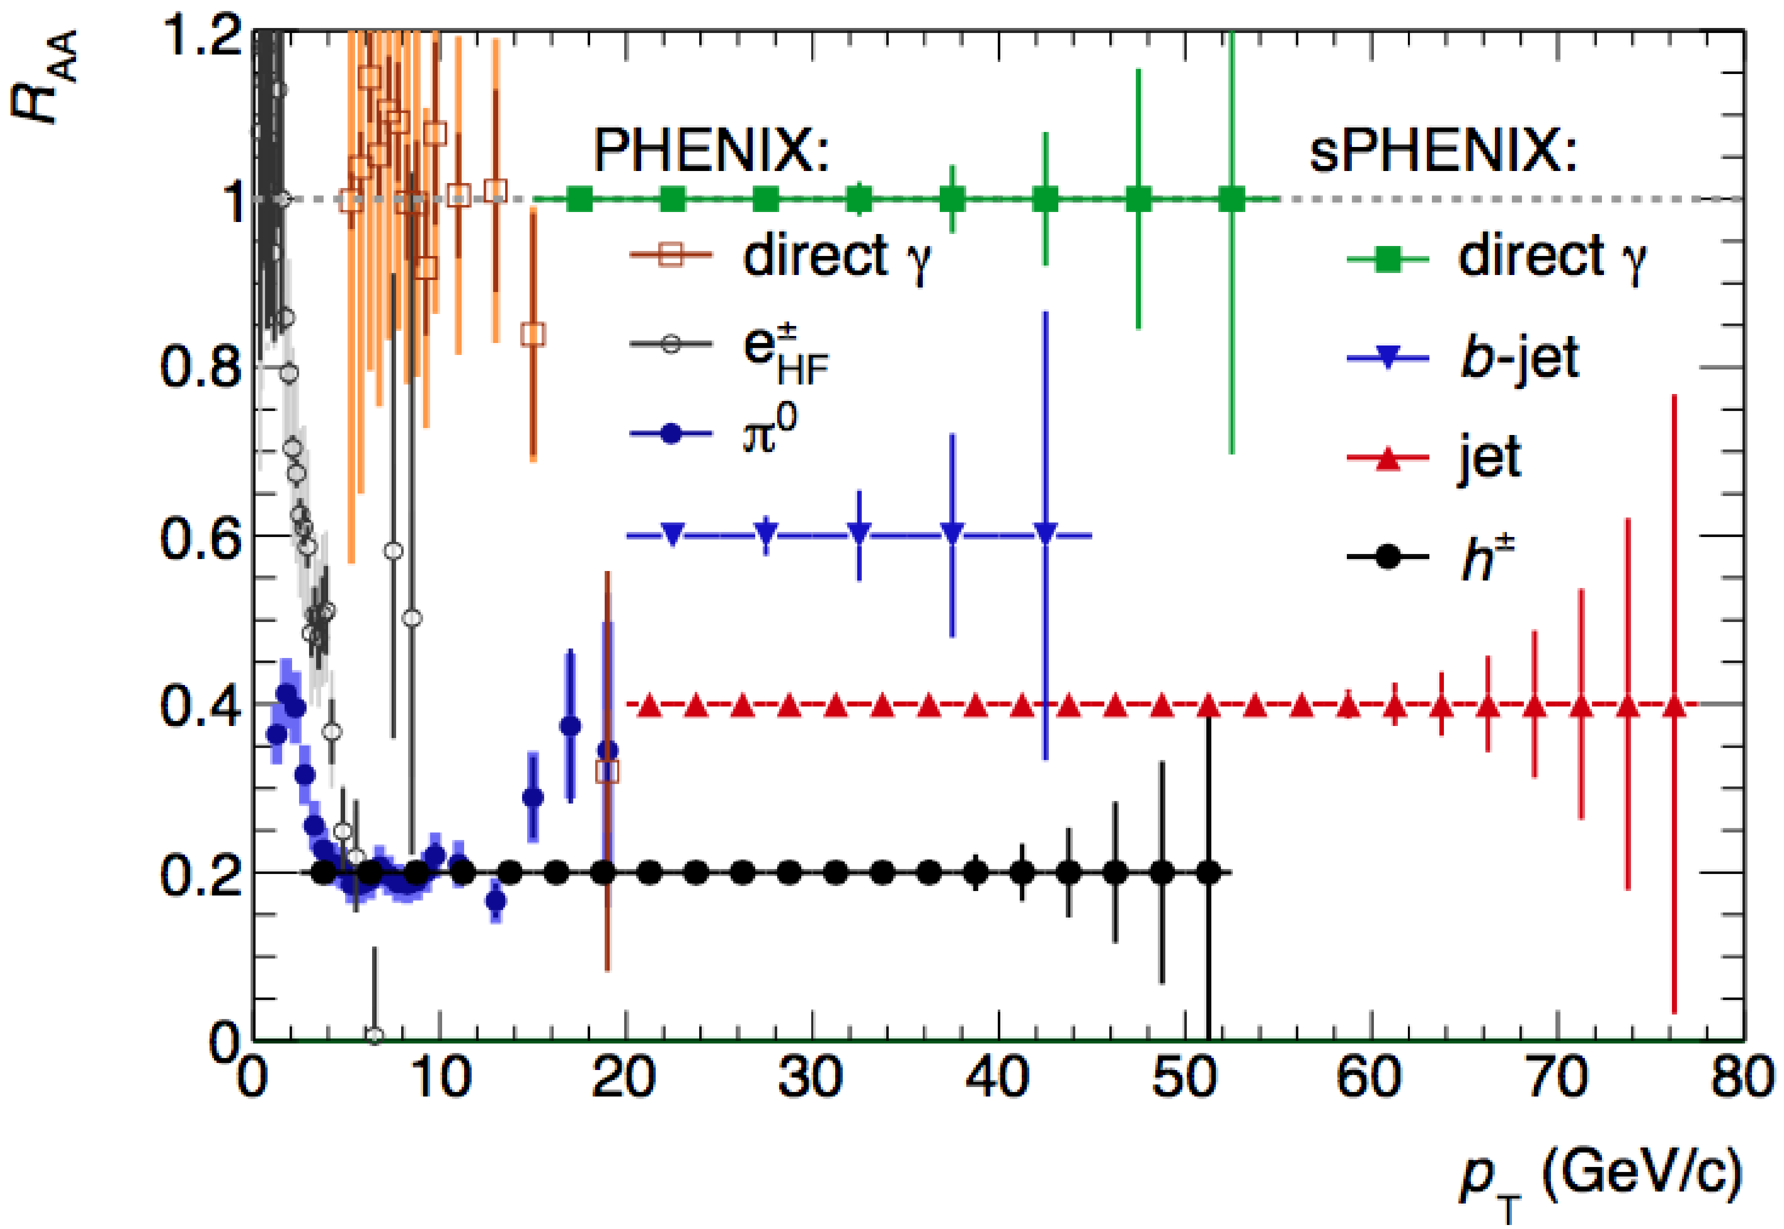

sPHENIX – the Collaboration

Upsilons

Just shy of 100 MeV/c resulution, enough to resolve the Upsilon states

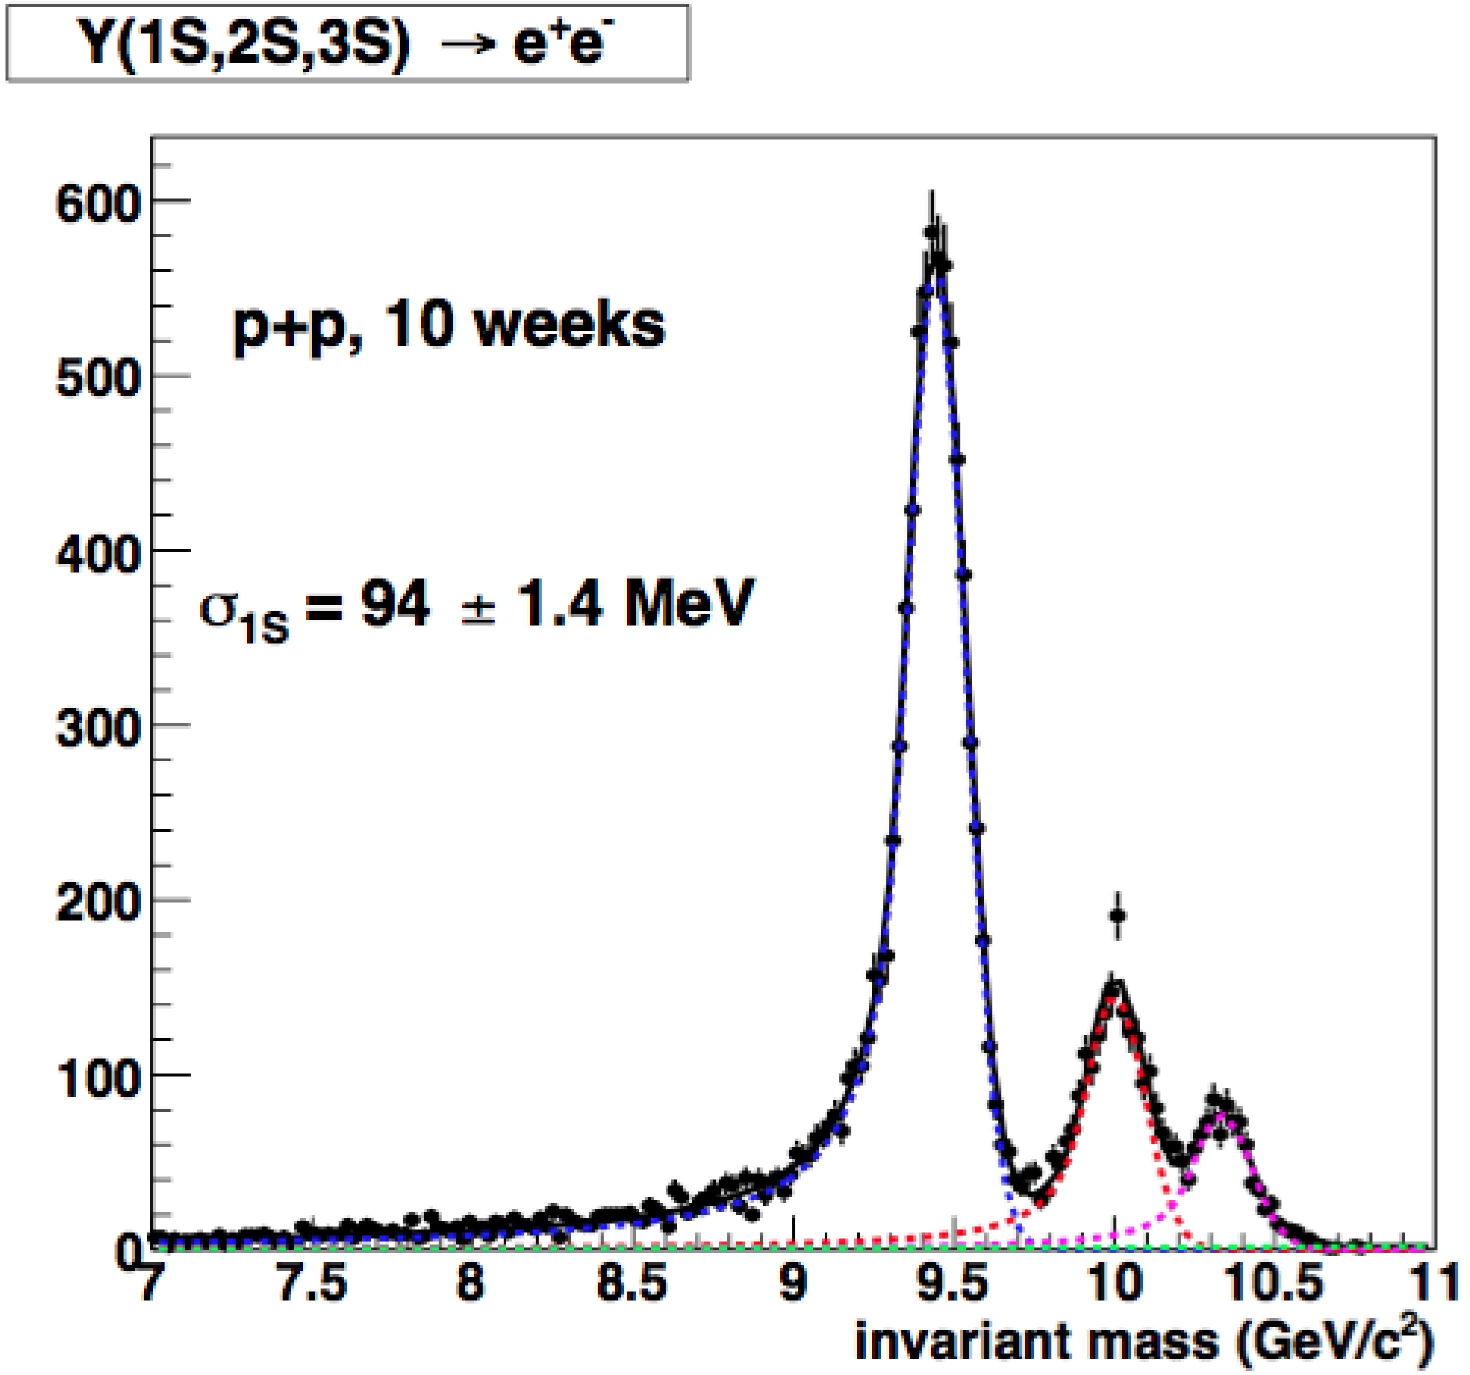

Tracking – 2 options, still very much under discussion

All Si Tracker option

Si tracker

- 7 layers strips and pixels
- Achieves design goals of pattern recognition and 100 MeV mass resolution on Upsilon states
- Total thickness $\approx 0.1X_0$

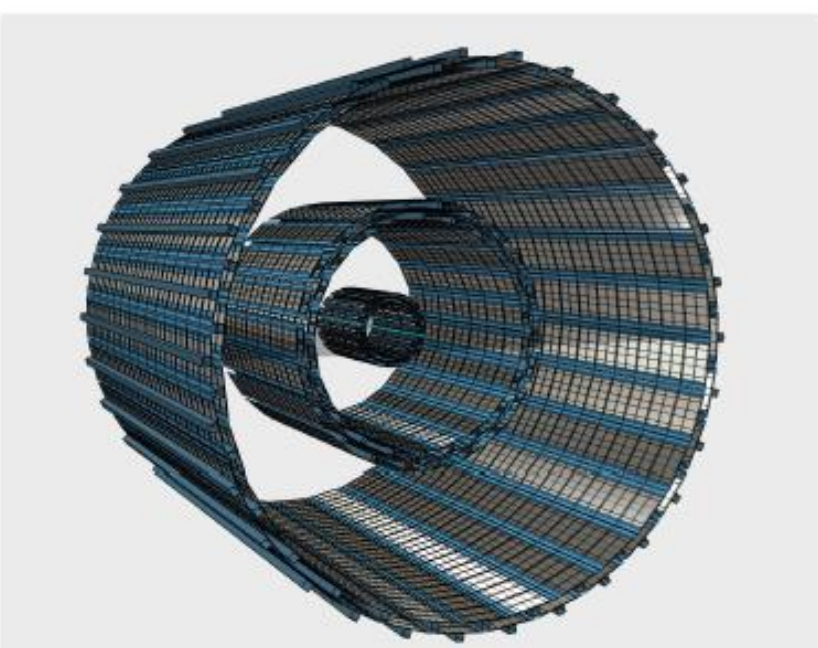

2 Pixel Layers + Compact TPC option

TPC + inner Si layers

- 80 cm outer radius TPC
- Inner Si detector
- TPC electronics following from ALICE upgrade

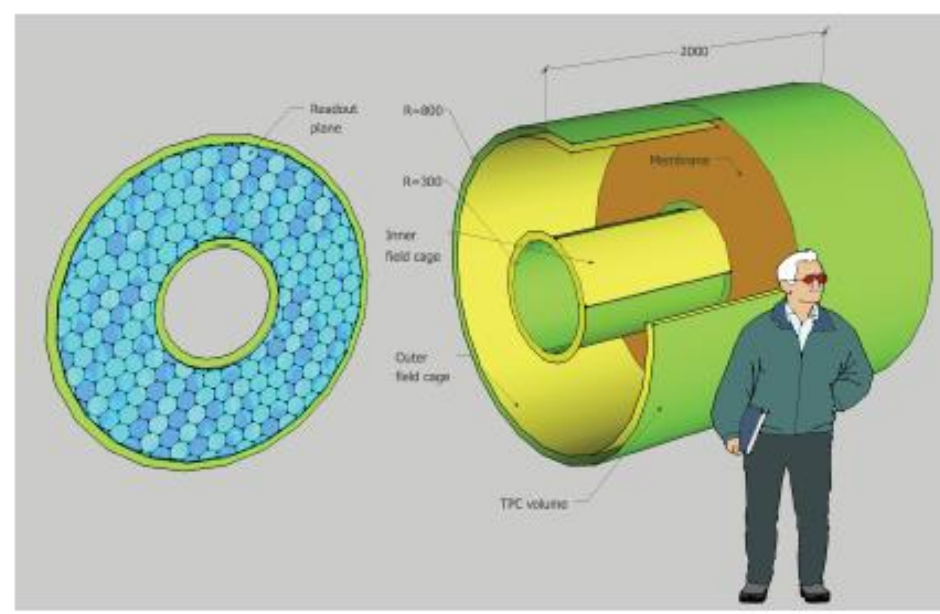

The BaBar Magnet

We had looked at magnets, but none of the available ones were quite right, then --

The BaBar magnet secured from SLAC after SuperB canceled, arrived at BNL in February 2015

Considerable additional equipment also acquired (power supplies, dump resistor, quench protection, cryogenic equipment)

SMD and CAD preparing it for low power cold test

Well suited to our needs without compromises

- 1.5 T central field
- 2.8 m diameter bore
- 3.8 m long
- $1.4X_0$ coil+cryostat

This really gave a tremendous boost to this project!

EmCal Modules

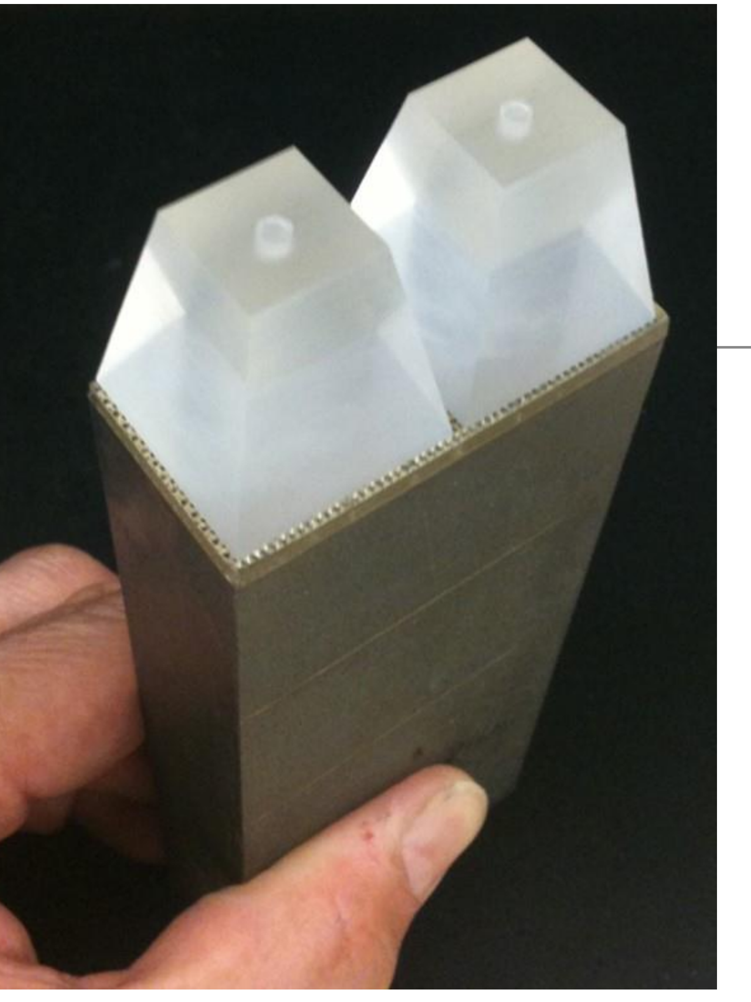

Tungsten-scintillating fiber SPACAL

Radiation length of ≈7 mm allows it to be inside the solenoid where only the material of the tracker is in front of it

Beam tested by UCLA group

Development of projective geometry which could improve e/π separation needed for the Upsilon measurements

Readout on inner radius of EMCAL with 4 3x3 mm SiPM's

On-detector electronics limited to preamps, bias control and temperature monitoring

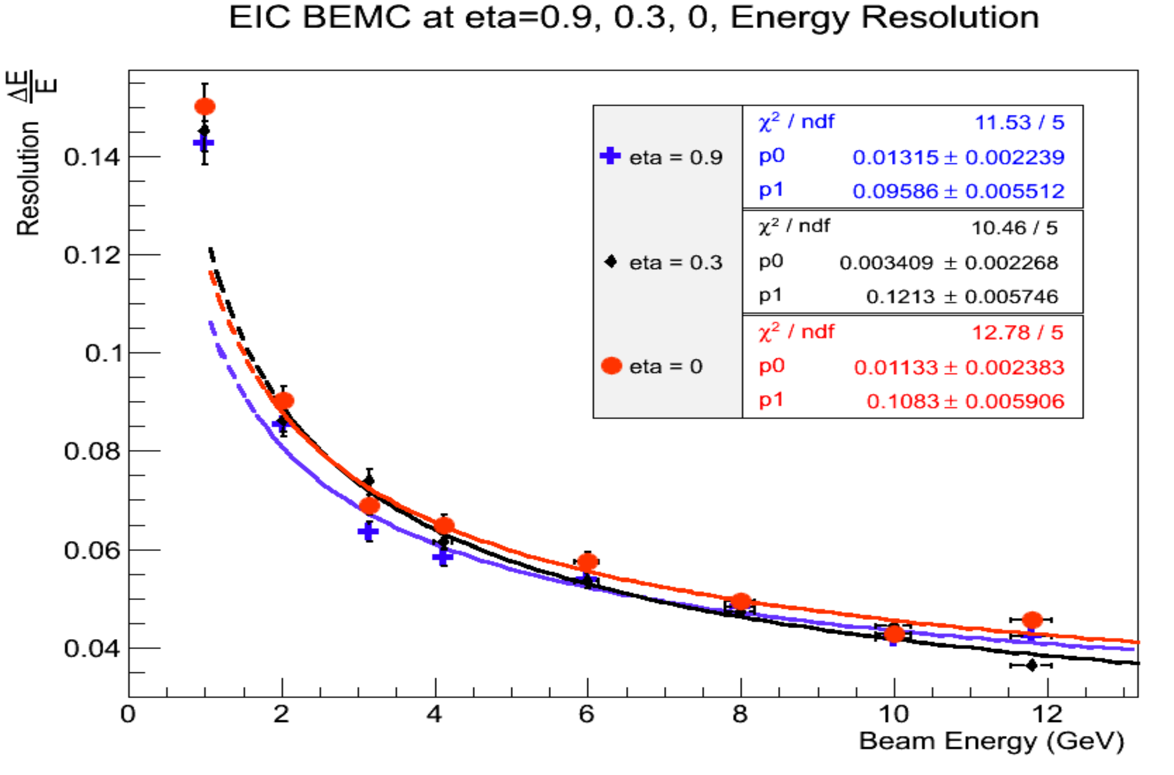

Projective Geometry

We want to make the EmCal projective, this is, each eta ring “looks” straight at the collision point

Superior over a “pineapple slice” design, angle of incidence is limited

That requires double-tapered modules to do it right

It is an engineering challenge though

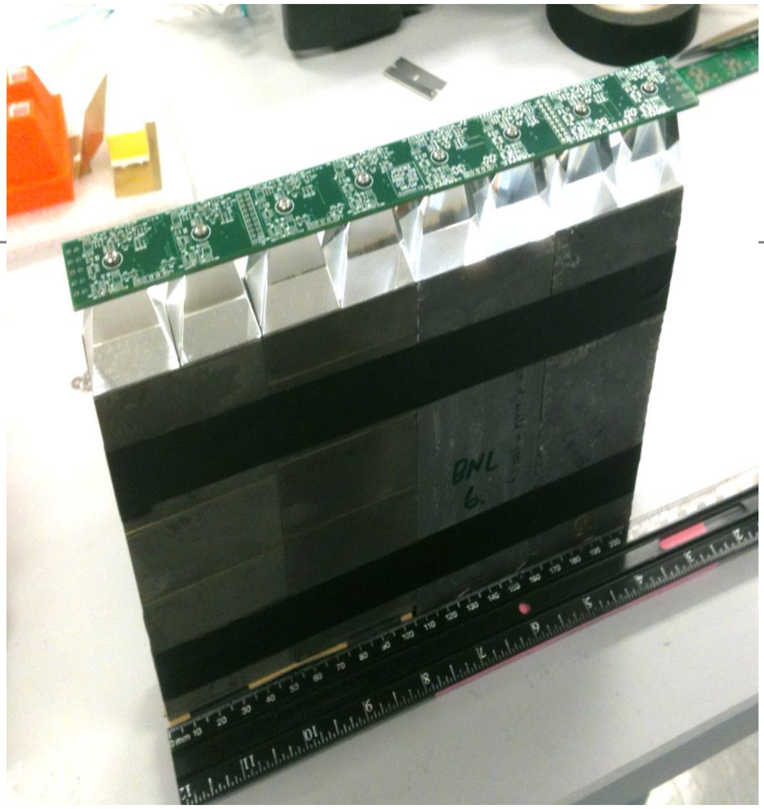

Ultimately want to build ~25k towers

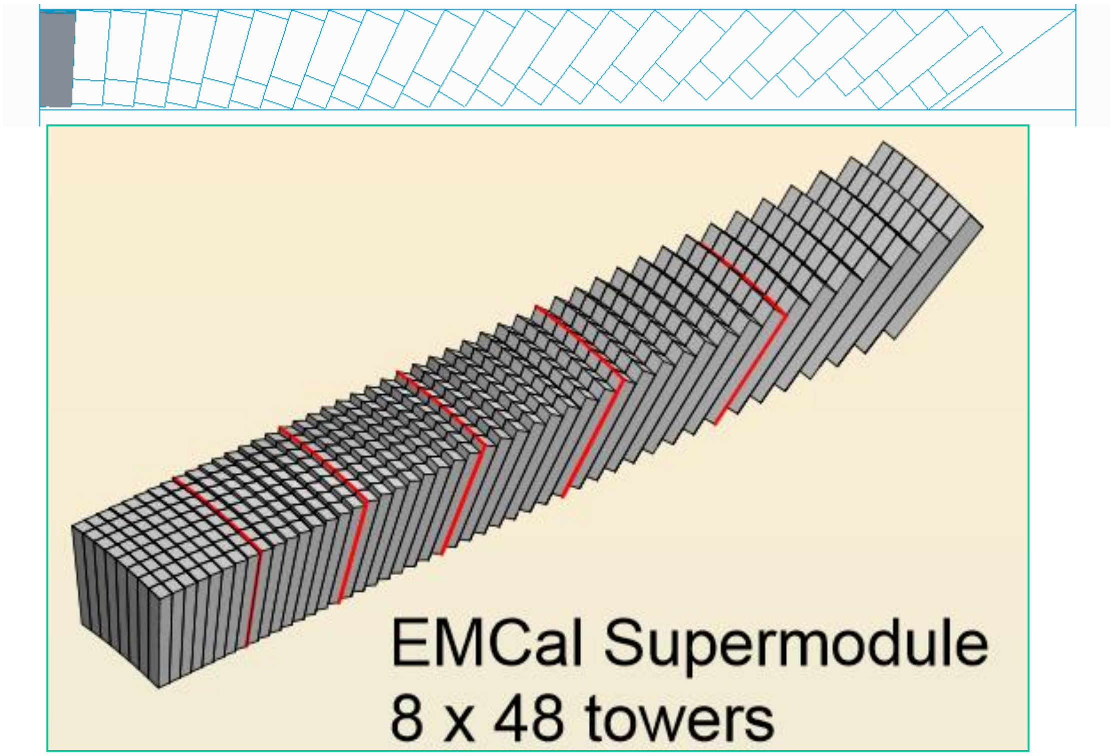

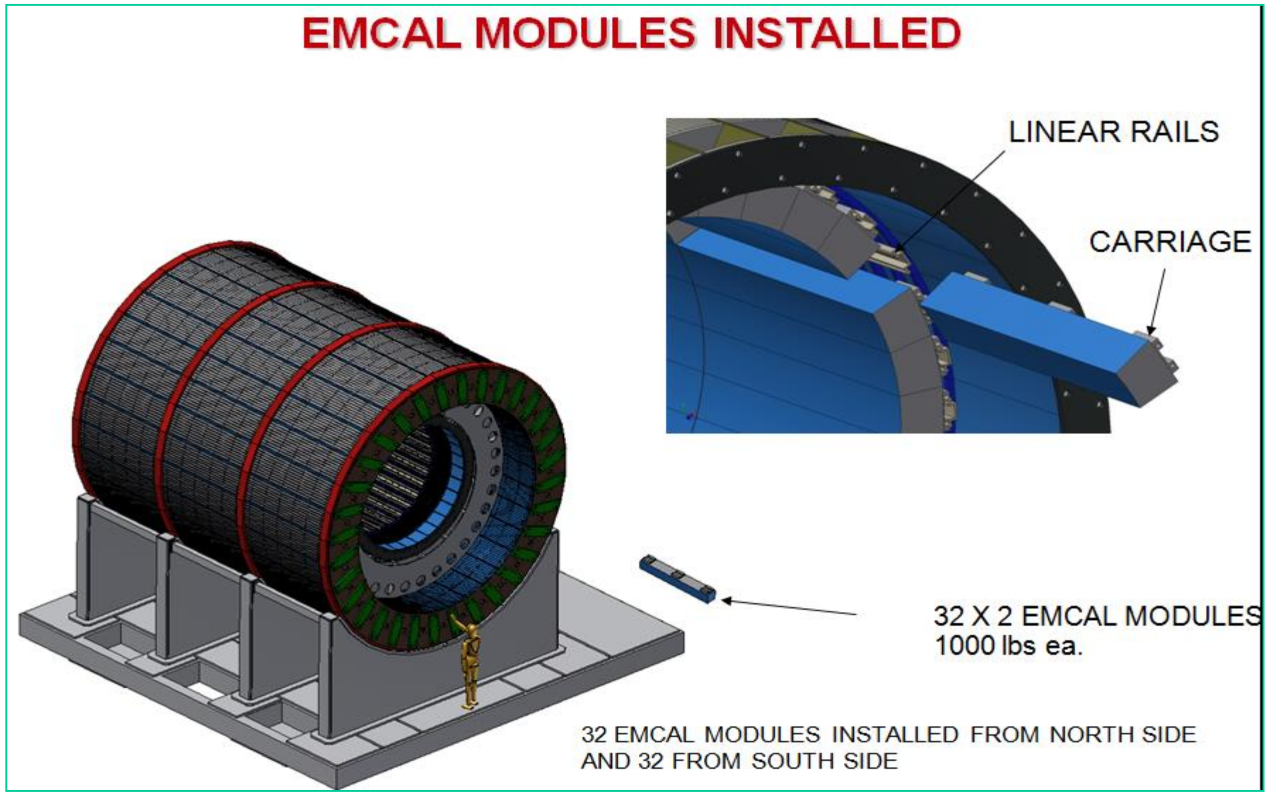

Density Uniformity, Sampling Fraction

- The tapered design makes the sampling fraction change over the depth of the module – we need to test the energy reconstruction carefully
- The production process is not “industrial-grade” yet, still at the level of prototypes
- Still a good amount of density variation module by module – we are still learning the production process.

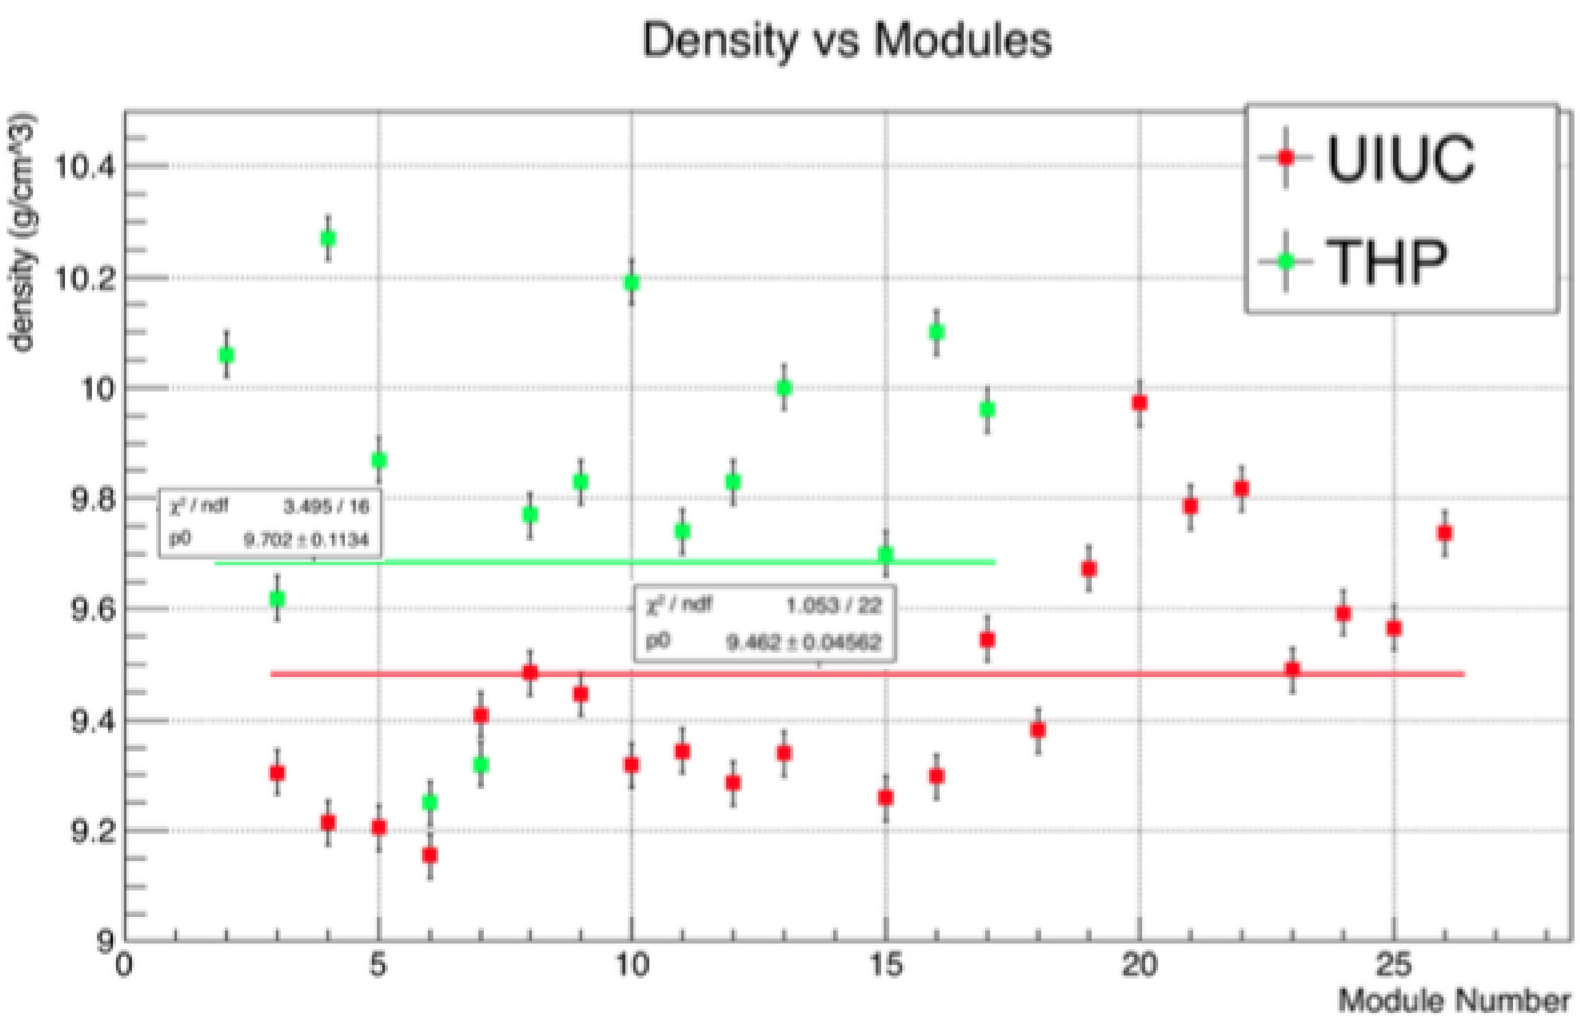

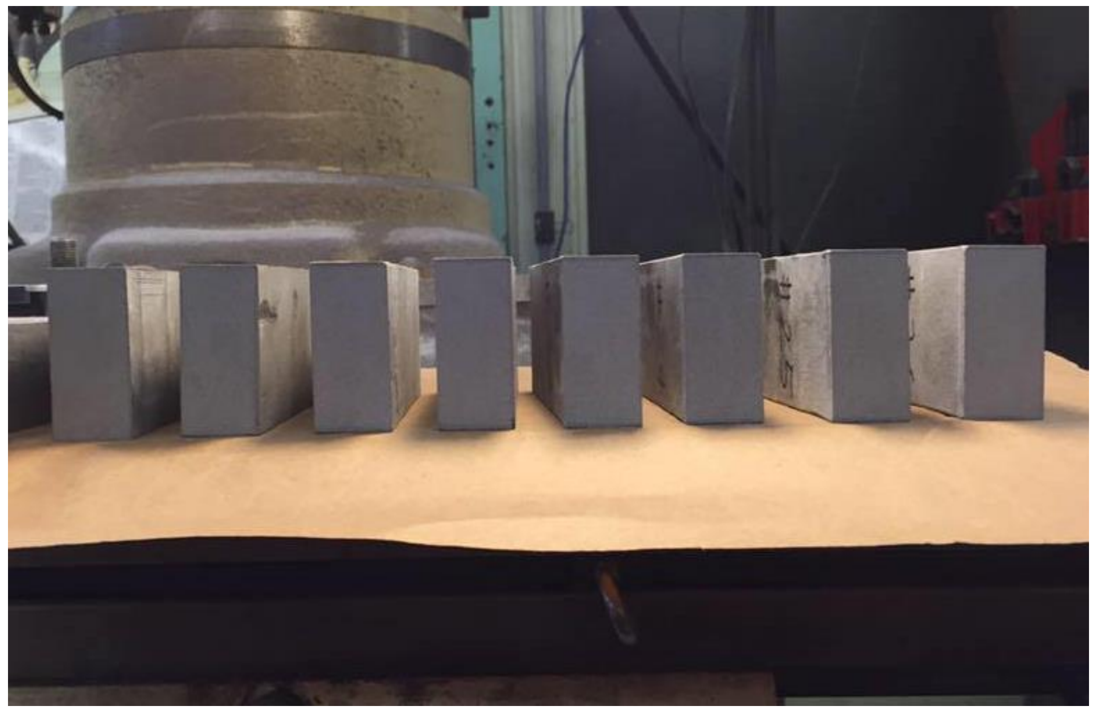

The latest batch of modules from Illinois

sPHENIX – step by step

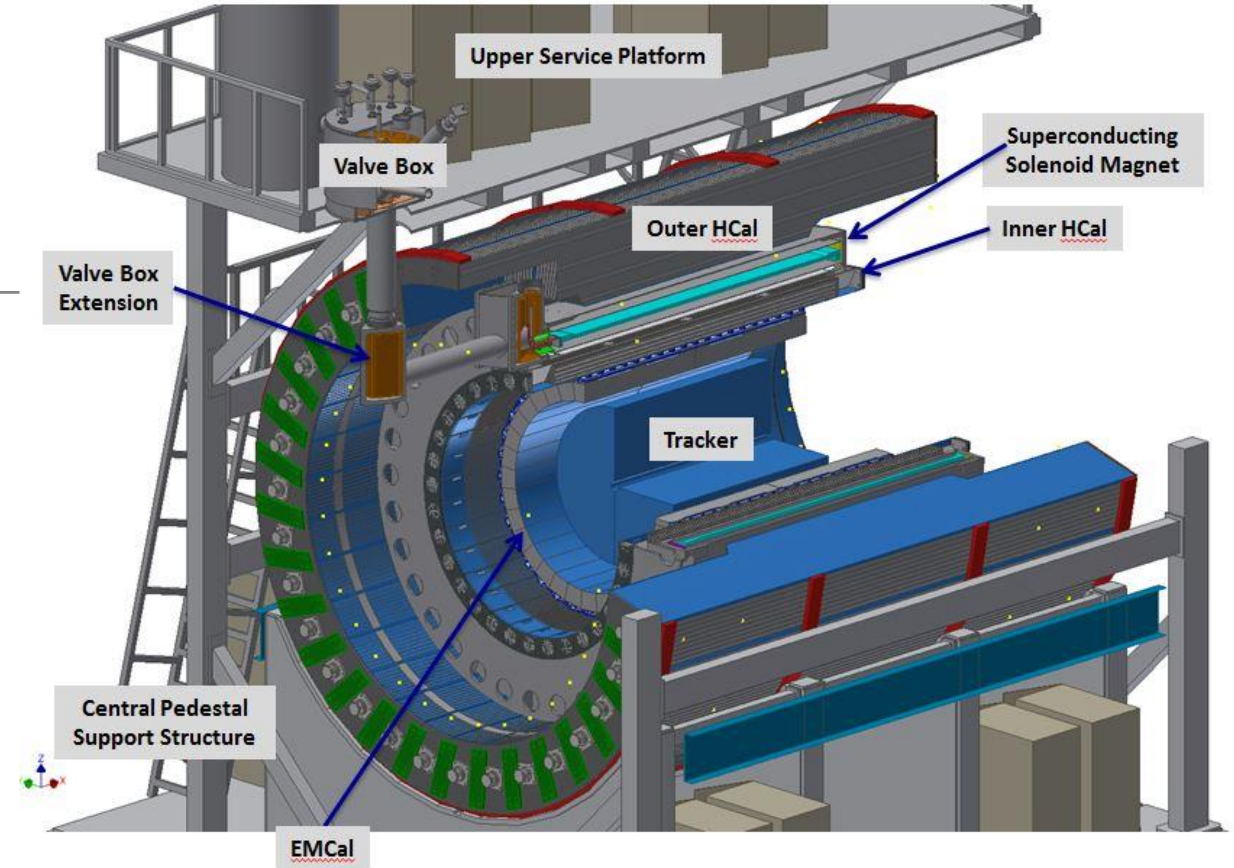

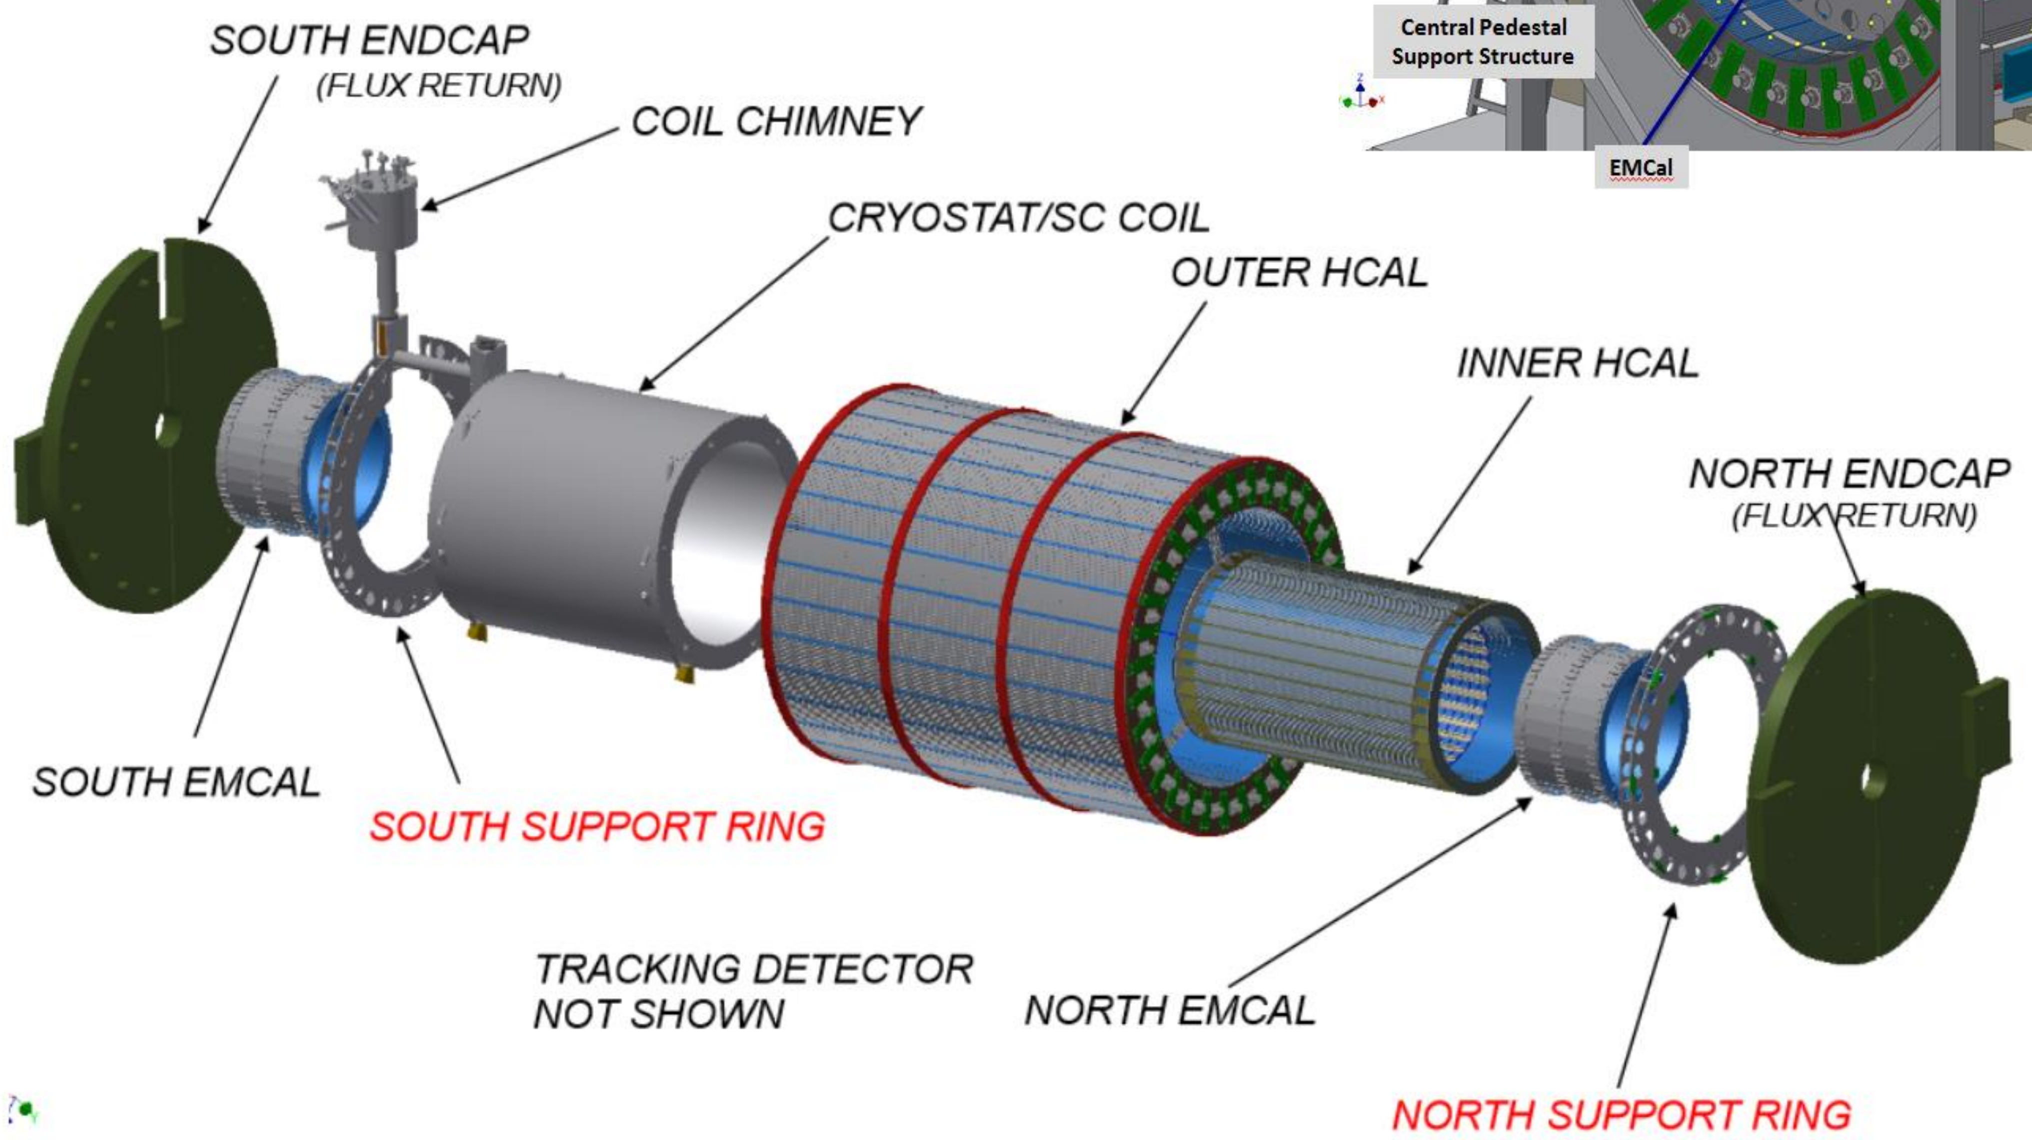

Inner HCal

32 sectors - 1.16m inner radius, 1.37m outer radius
10 rows of 7mm scintillator tiles (24 tiles per row)
32^{o} tilt angle, tapered stainless steel plates ~10.2mm - ~14.7mm
Completed sector is 4.3m long, weighs ~ 1 ton

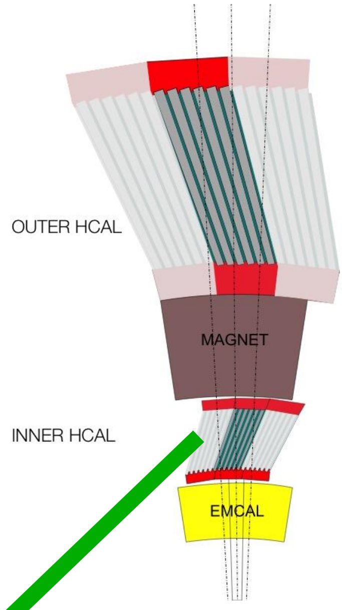

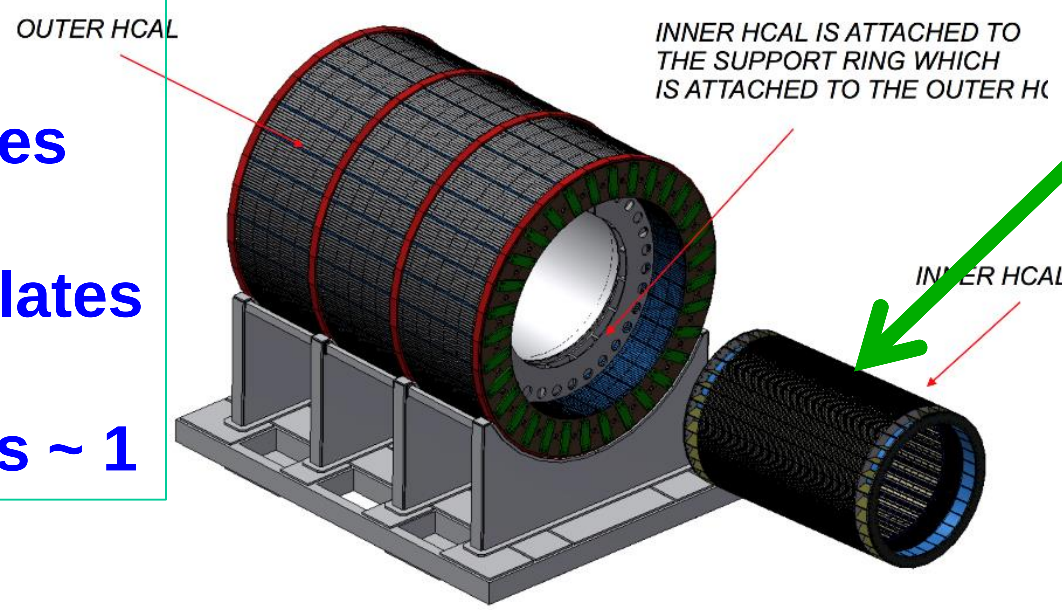

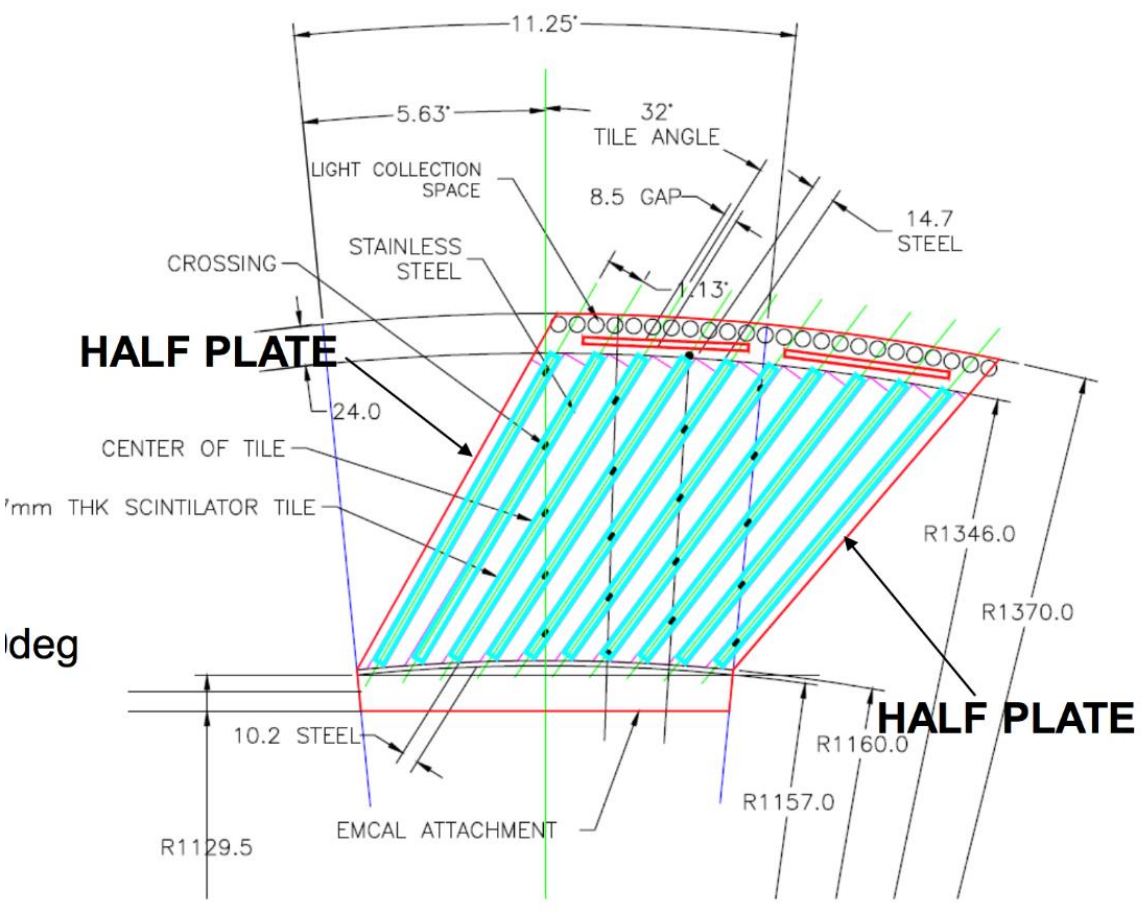

24

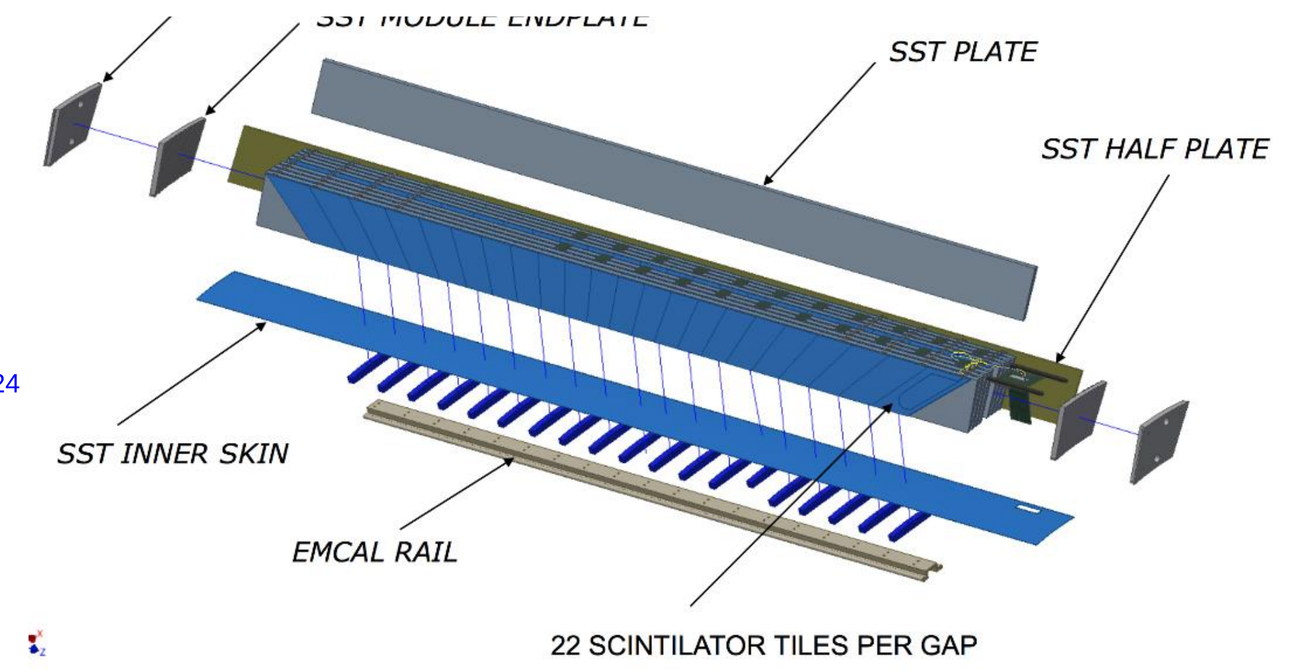

Outer HCal

32 sectors - 1.9m inner radius, 2.6m outer radius

10 rows of 7mm scint. tiles (24 tiles per row), 12^o tilt angle

Tapered 1006 steel plates ~26.1mm - ~42.4mm

6.3 m long, 13.5 tons

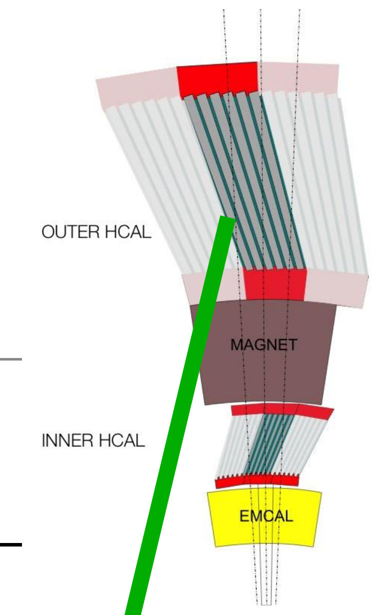

The drawing looks at the Hcal from "the back"

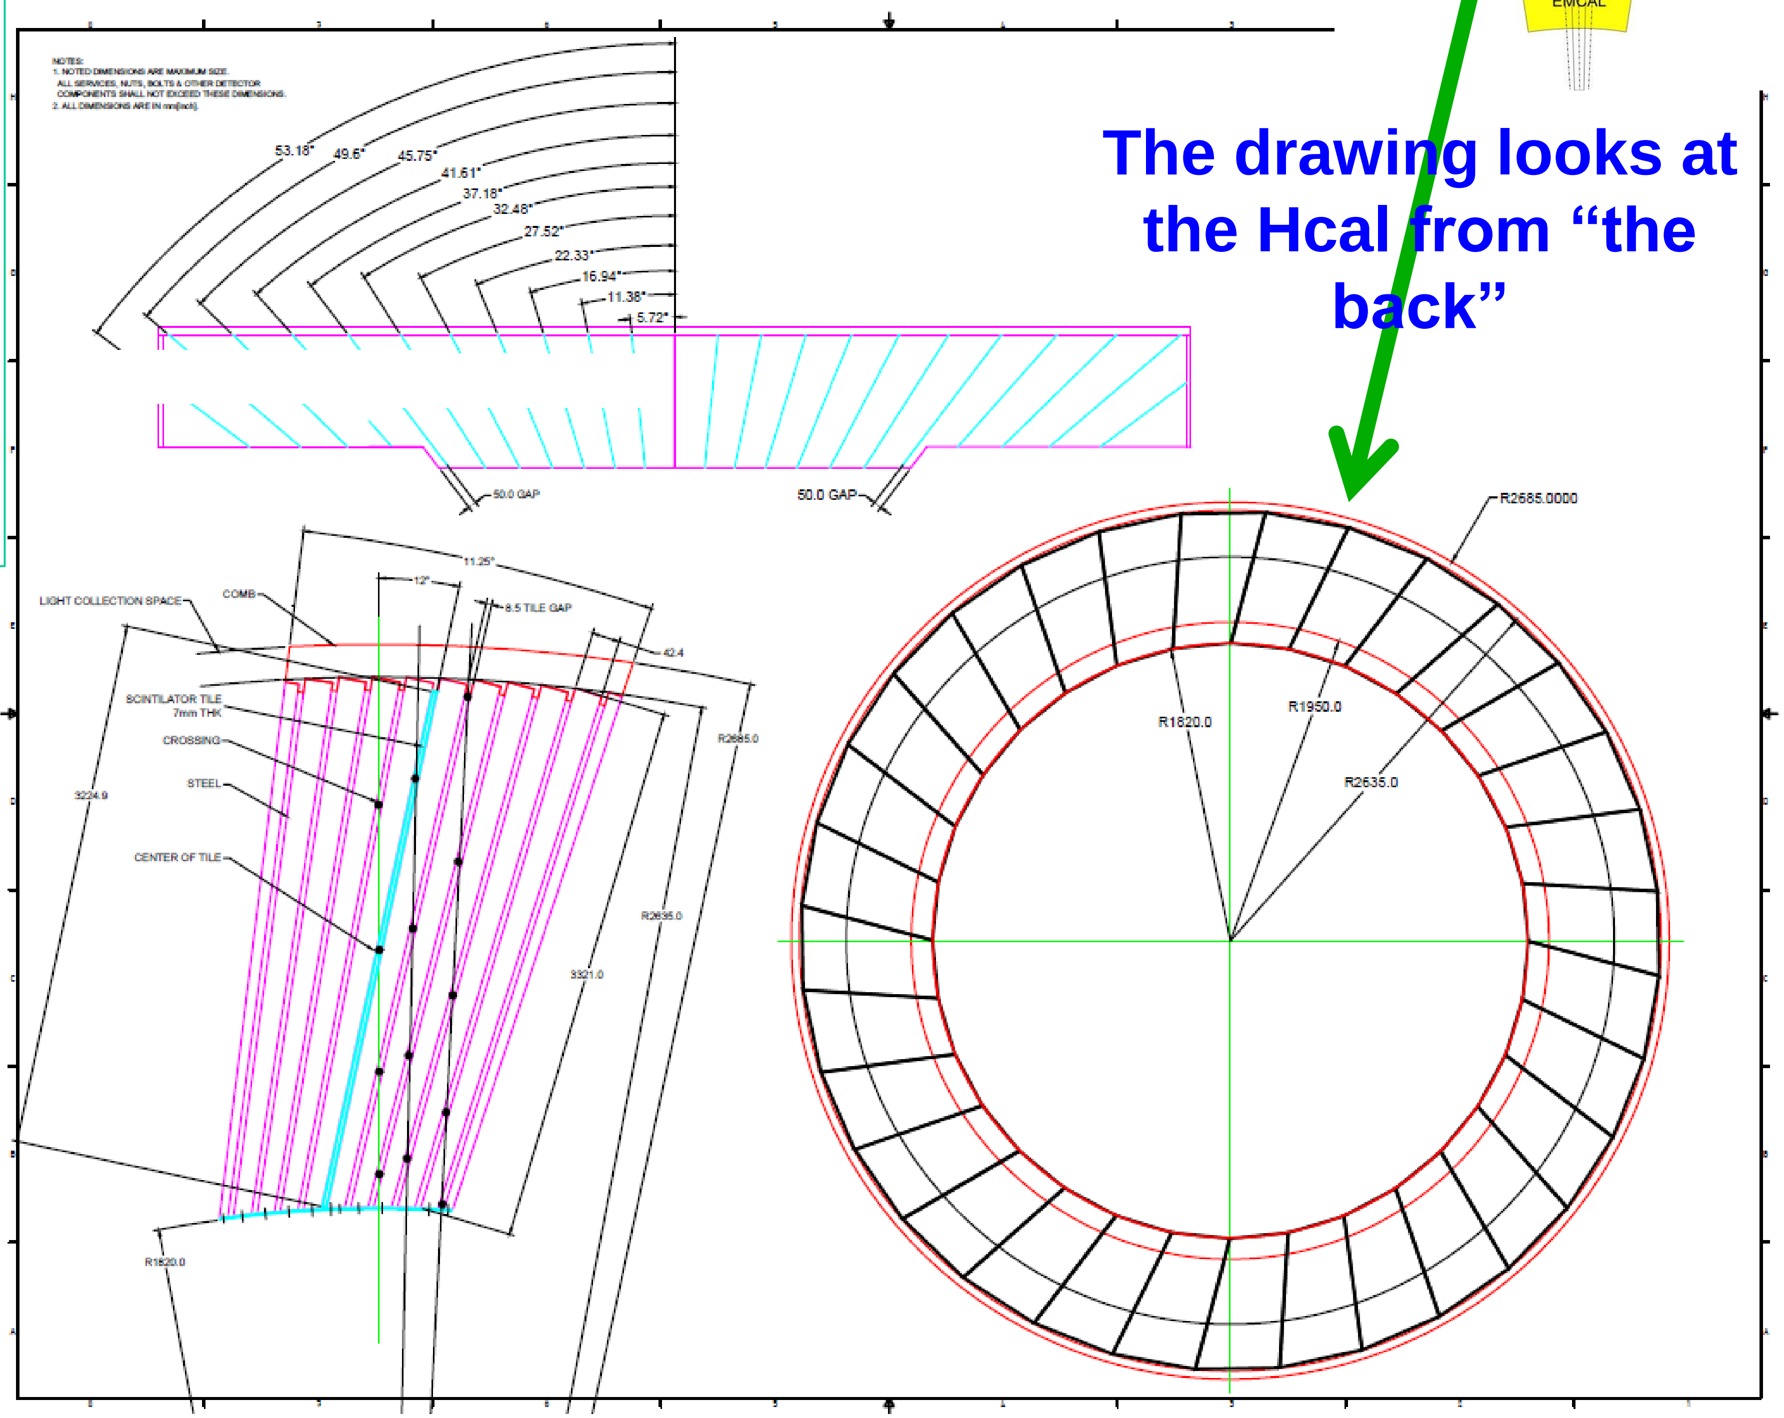

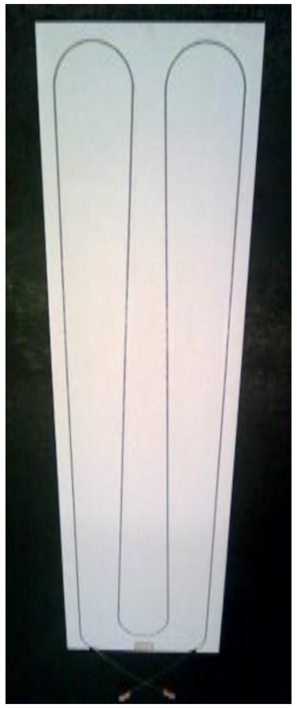

Scintillator tile

Just because of the weight, the outer HCal tends to dominate our engineering discussions!

What we stick into the Test Beams…

A “sector” of the HCals, plus a small EmCal portion

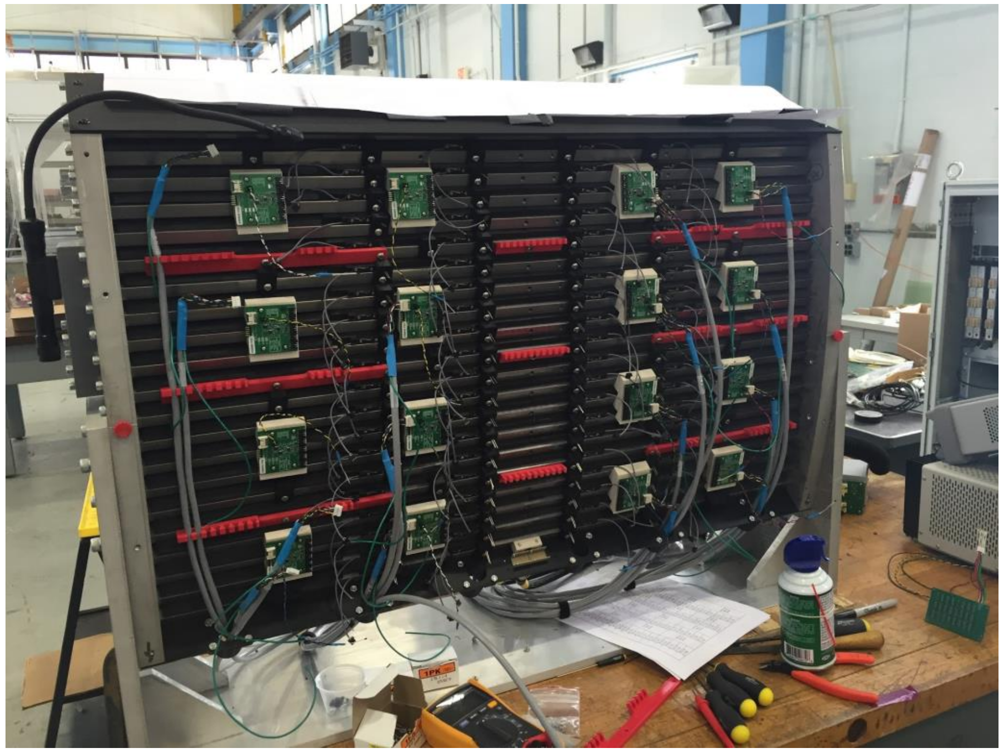

Multi-Event buffering

- This is what makes the DAQ as fast as it is
- Or, in other words: that keeps the dead time in the low 90%'s
- Comes down to dealing with your data from various events in parallel

An another-world example

A Volkswagen assembly line

A given car takes about 28 hours from starting as a naked chassis to being an assembled vehicle

One station adds the “skin”, another the engine, another installs the defeat devices – about 340 “stations”

- Single-Event buffering – next car enters the assembly line once the previous car is done → **one car every 28 hours**
- Multi-Event Buffering – car moves forward as soon as the next station is free → **one car about every 5 minutes**

Multi-Event buffering effect in the DAQ

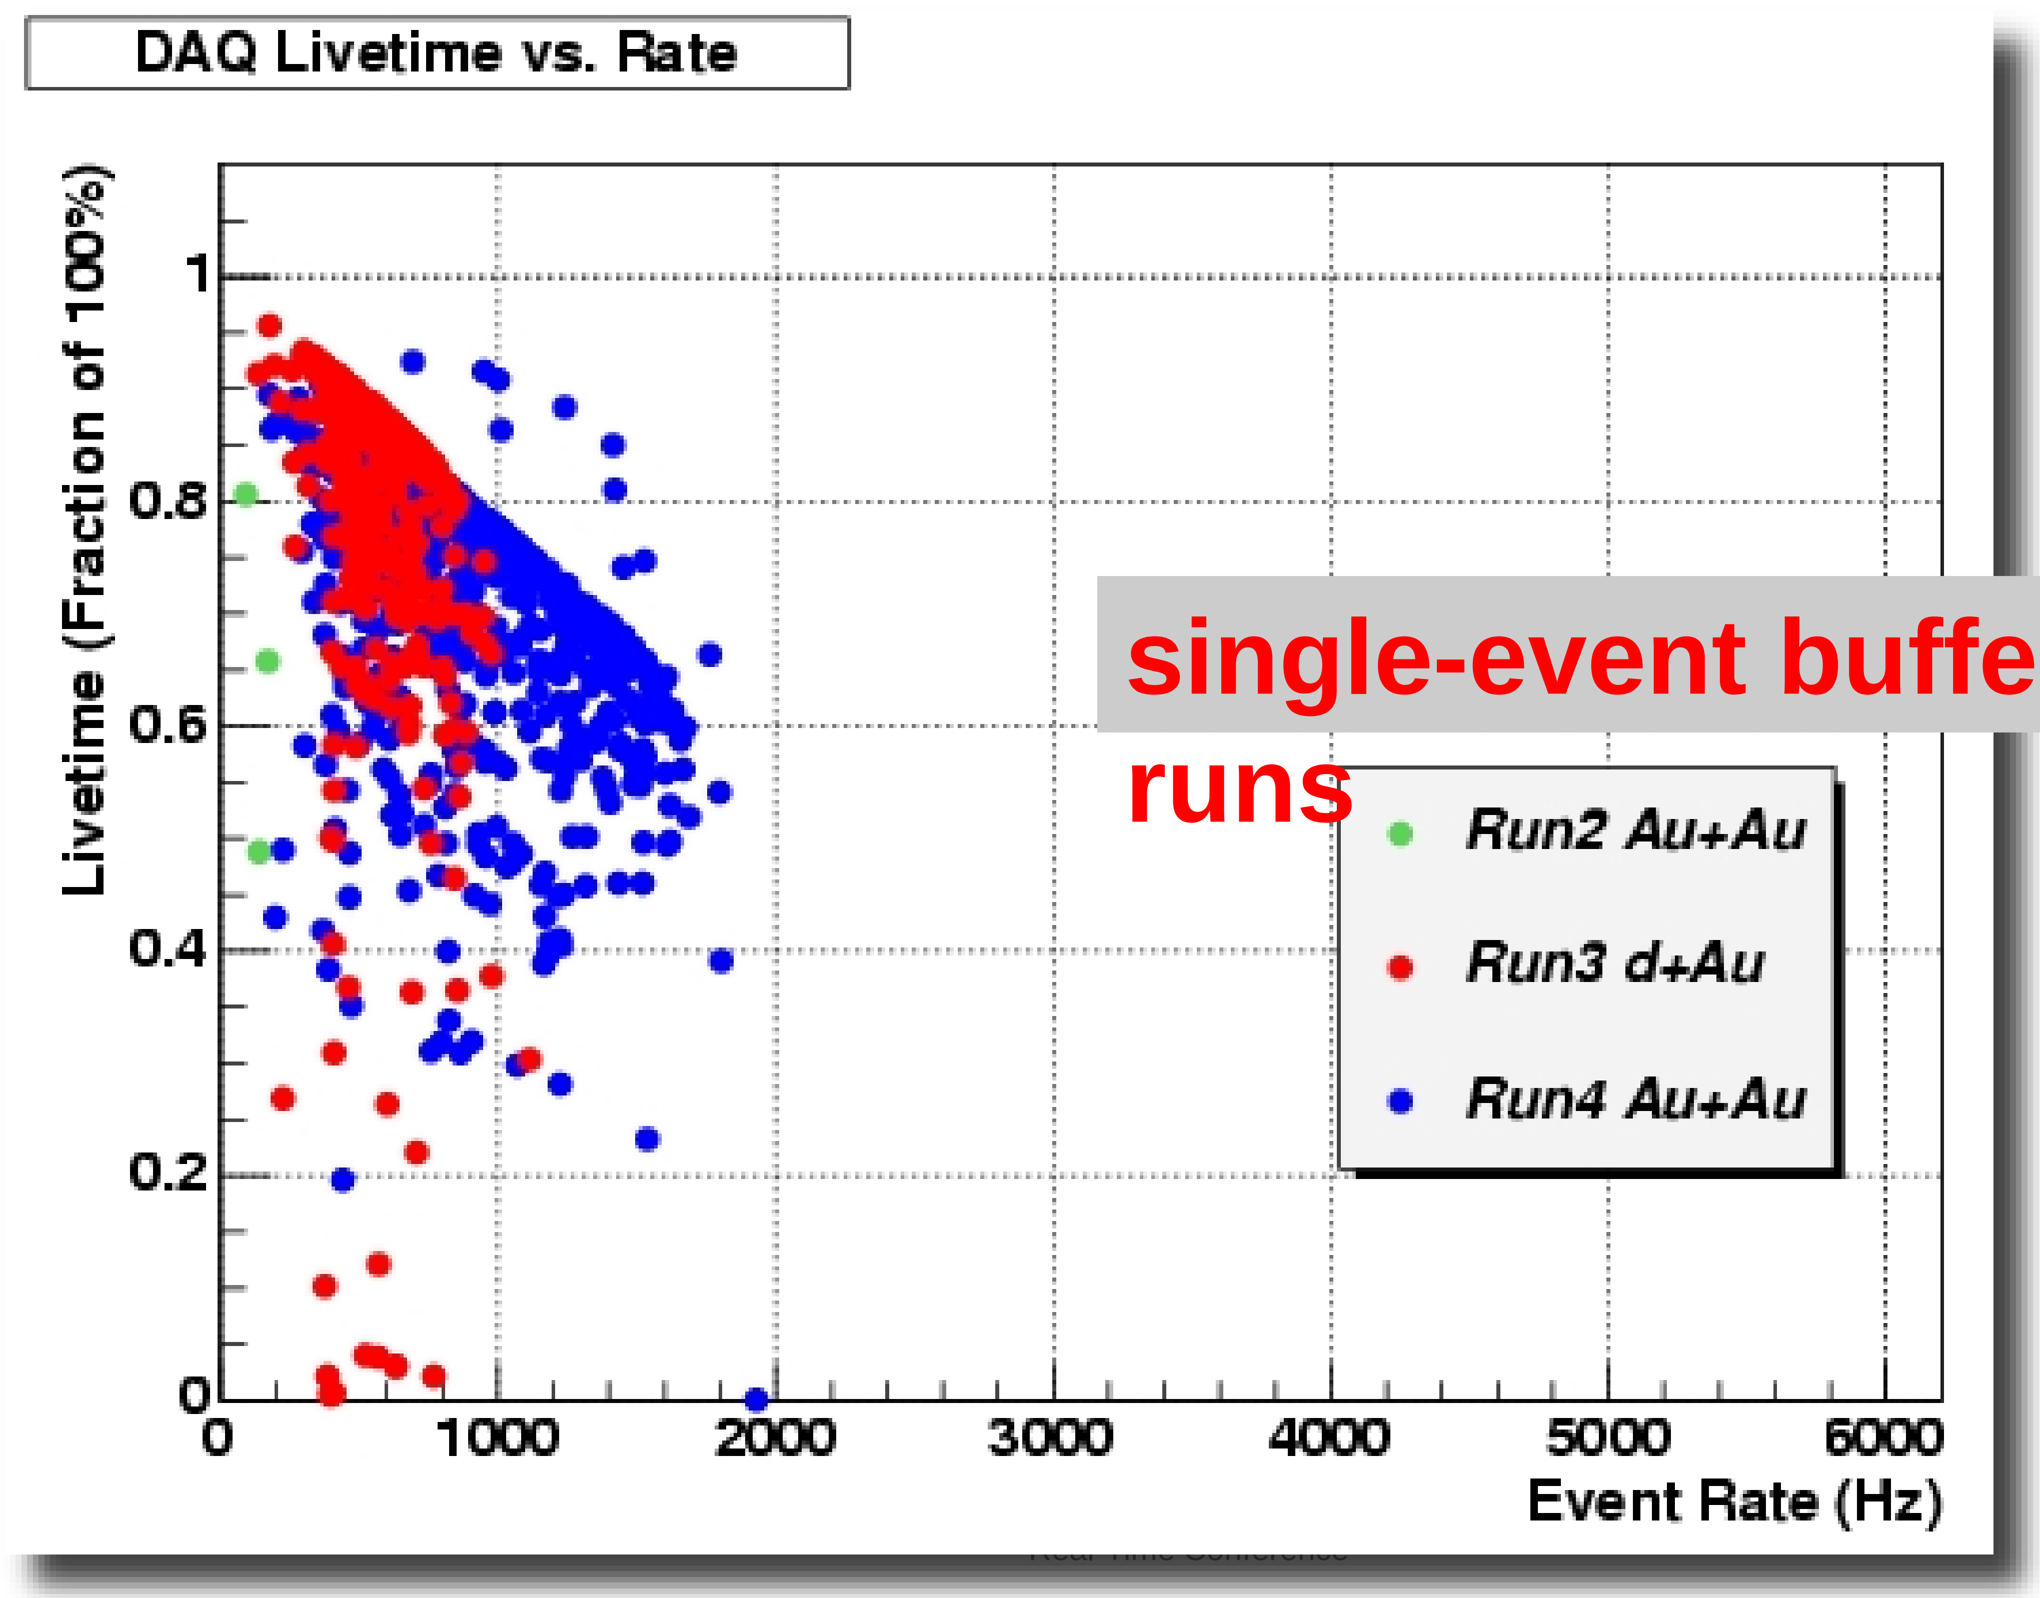

Multi-Event buffering effect in the DAQ

Multi-Event buffering effect in the DAQ

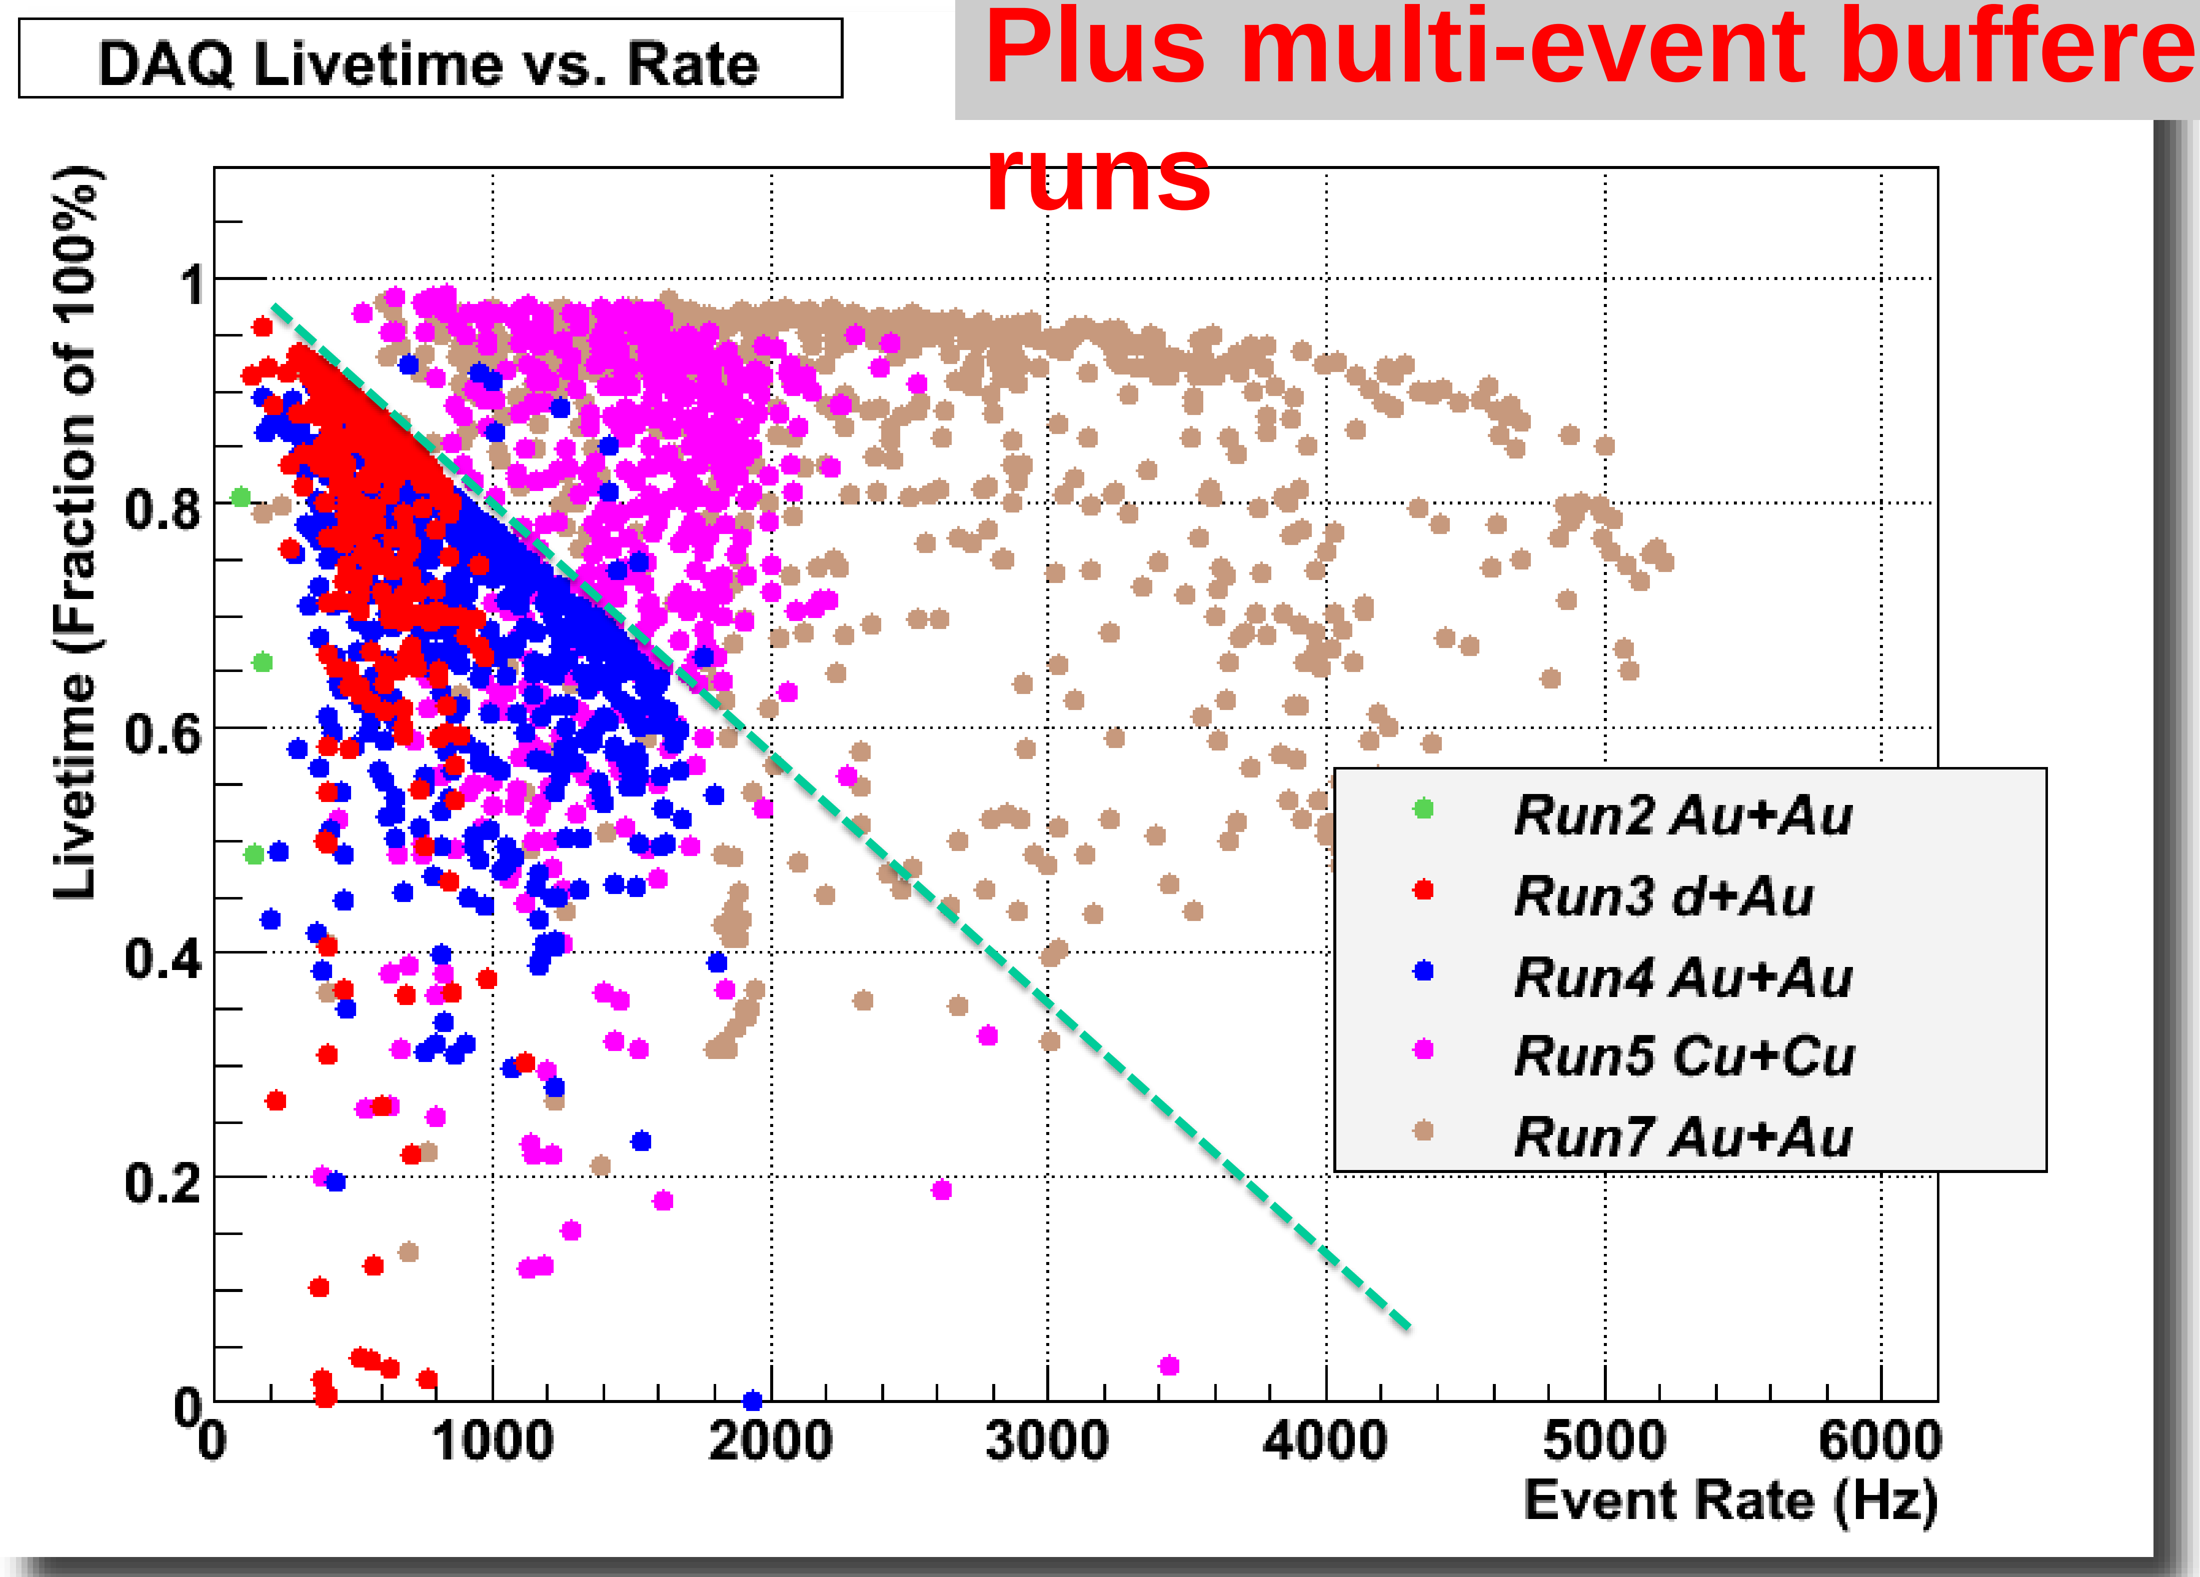

Plus multi-event buffered runs