DAQ Status and Computing Needs

Matt Graham, deputy L3 manager
SLAC
S&C Face-to-Face @ Fermilab
November 13, 2013

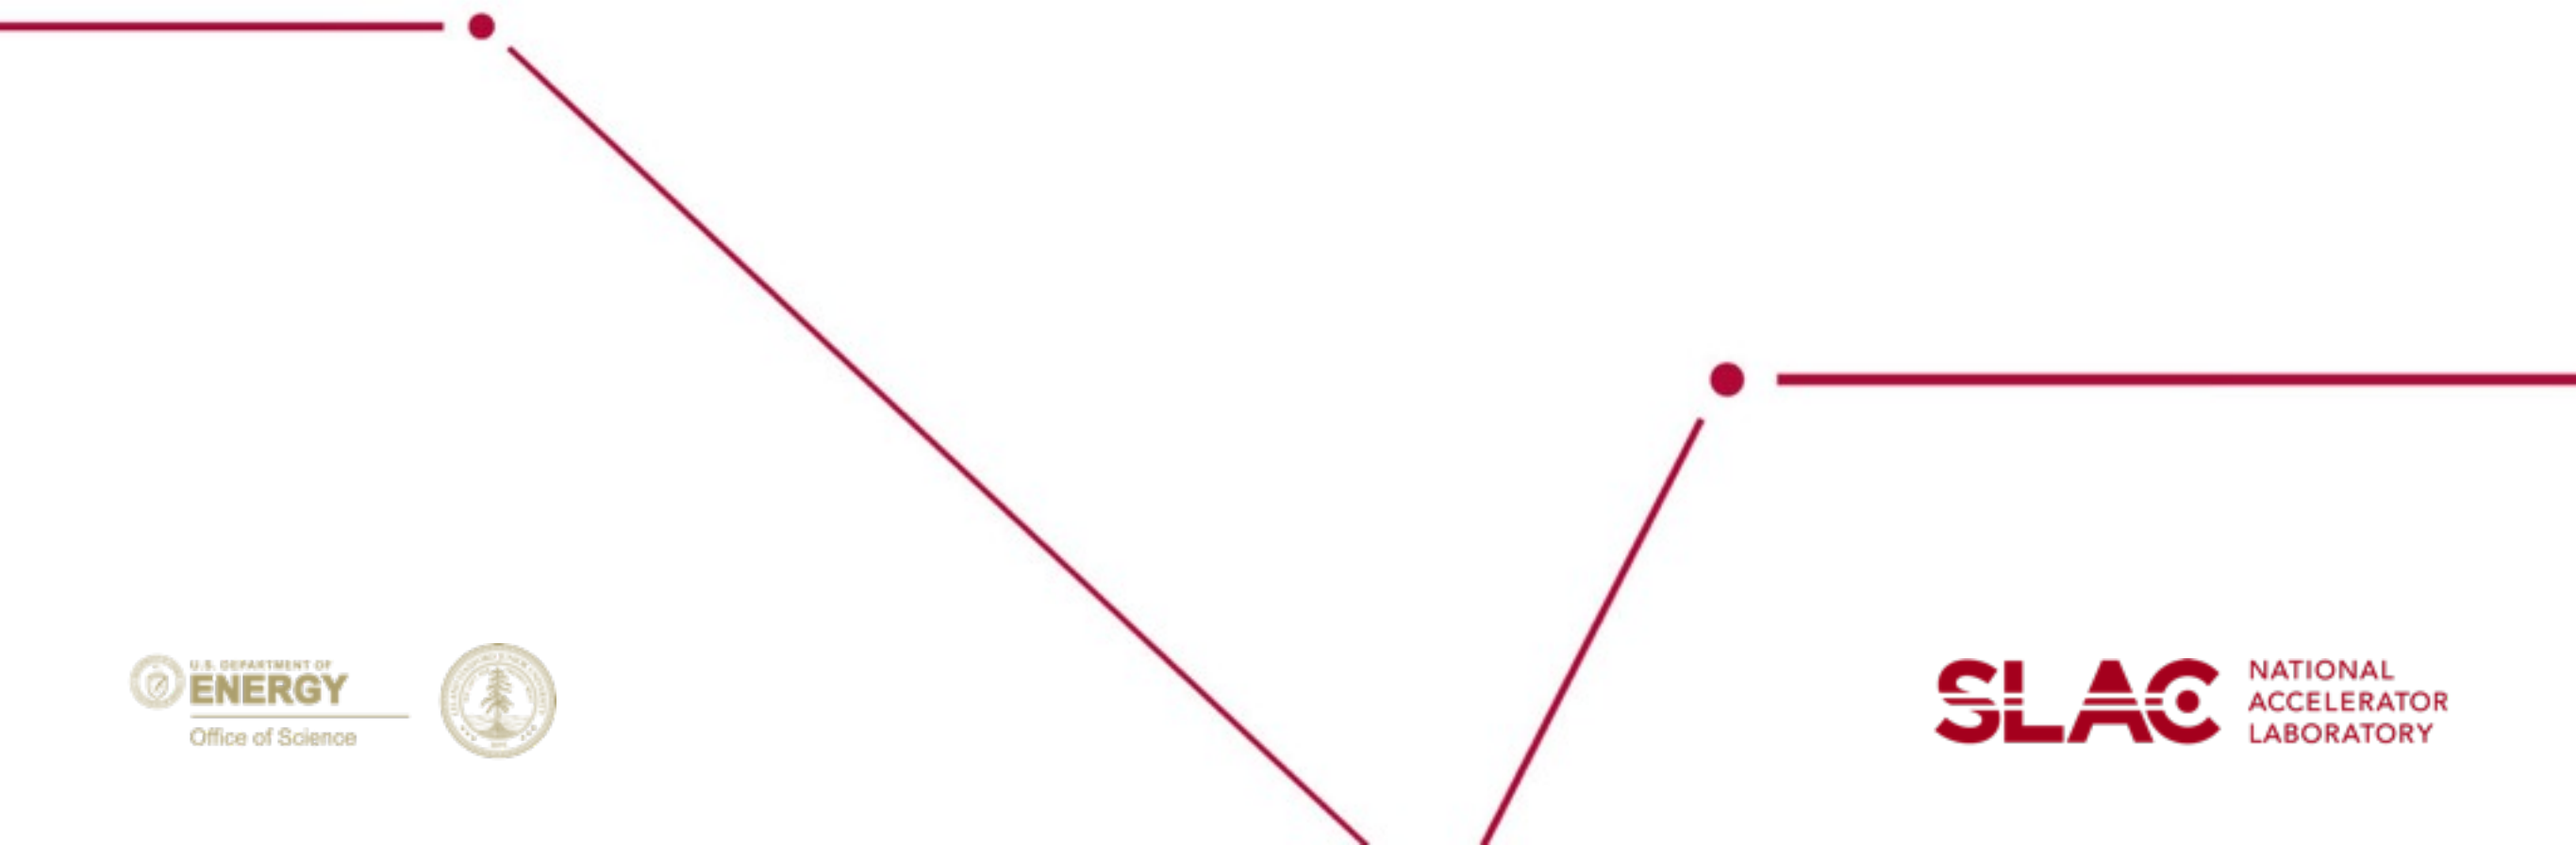

Overview

- Our current focus is on getting a working DAQ system for 35T phase 2.
  - We'll learn a lot from the prototype, but I've tried to put together our current thinking for full detector too.
- Two systems (for full detector
  - TPC : ~ 128 channels/FEB x 20 FEB/APA =2560 channels/APA (how many APAs?...was 64 for 17kt)
  - PD : ~ 10 Paddles/APA* 12 channels/Paddle=120 channels/APA
- For this talk, I'll go through our current thinking...some of this is pretty set, some less so.
- dataflow document docdb-7450 has a lot of detail for the 35t
  - is a bit out of date now, but I'll try to update it soon (I've had this same sentence in here for ~ 3 months now)
- For you guys, what's important to us: artDAQ, database management, backend DAQ computing (online farm)

DAQ High-level Block Diagram For Full Detector

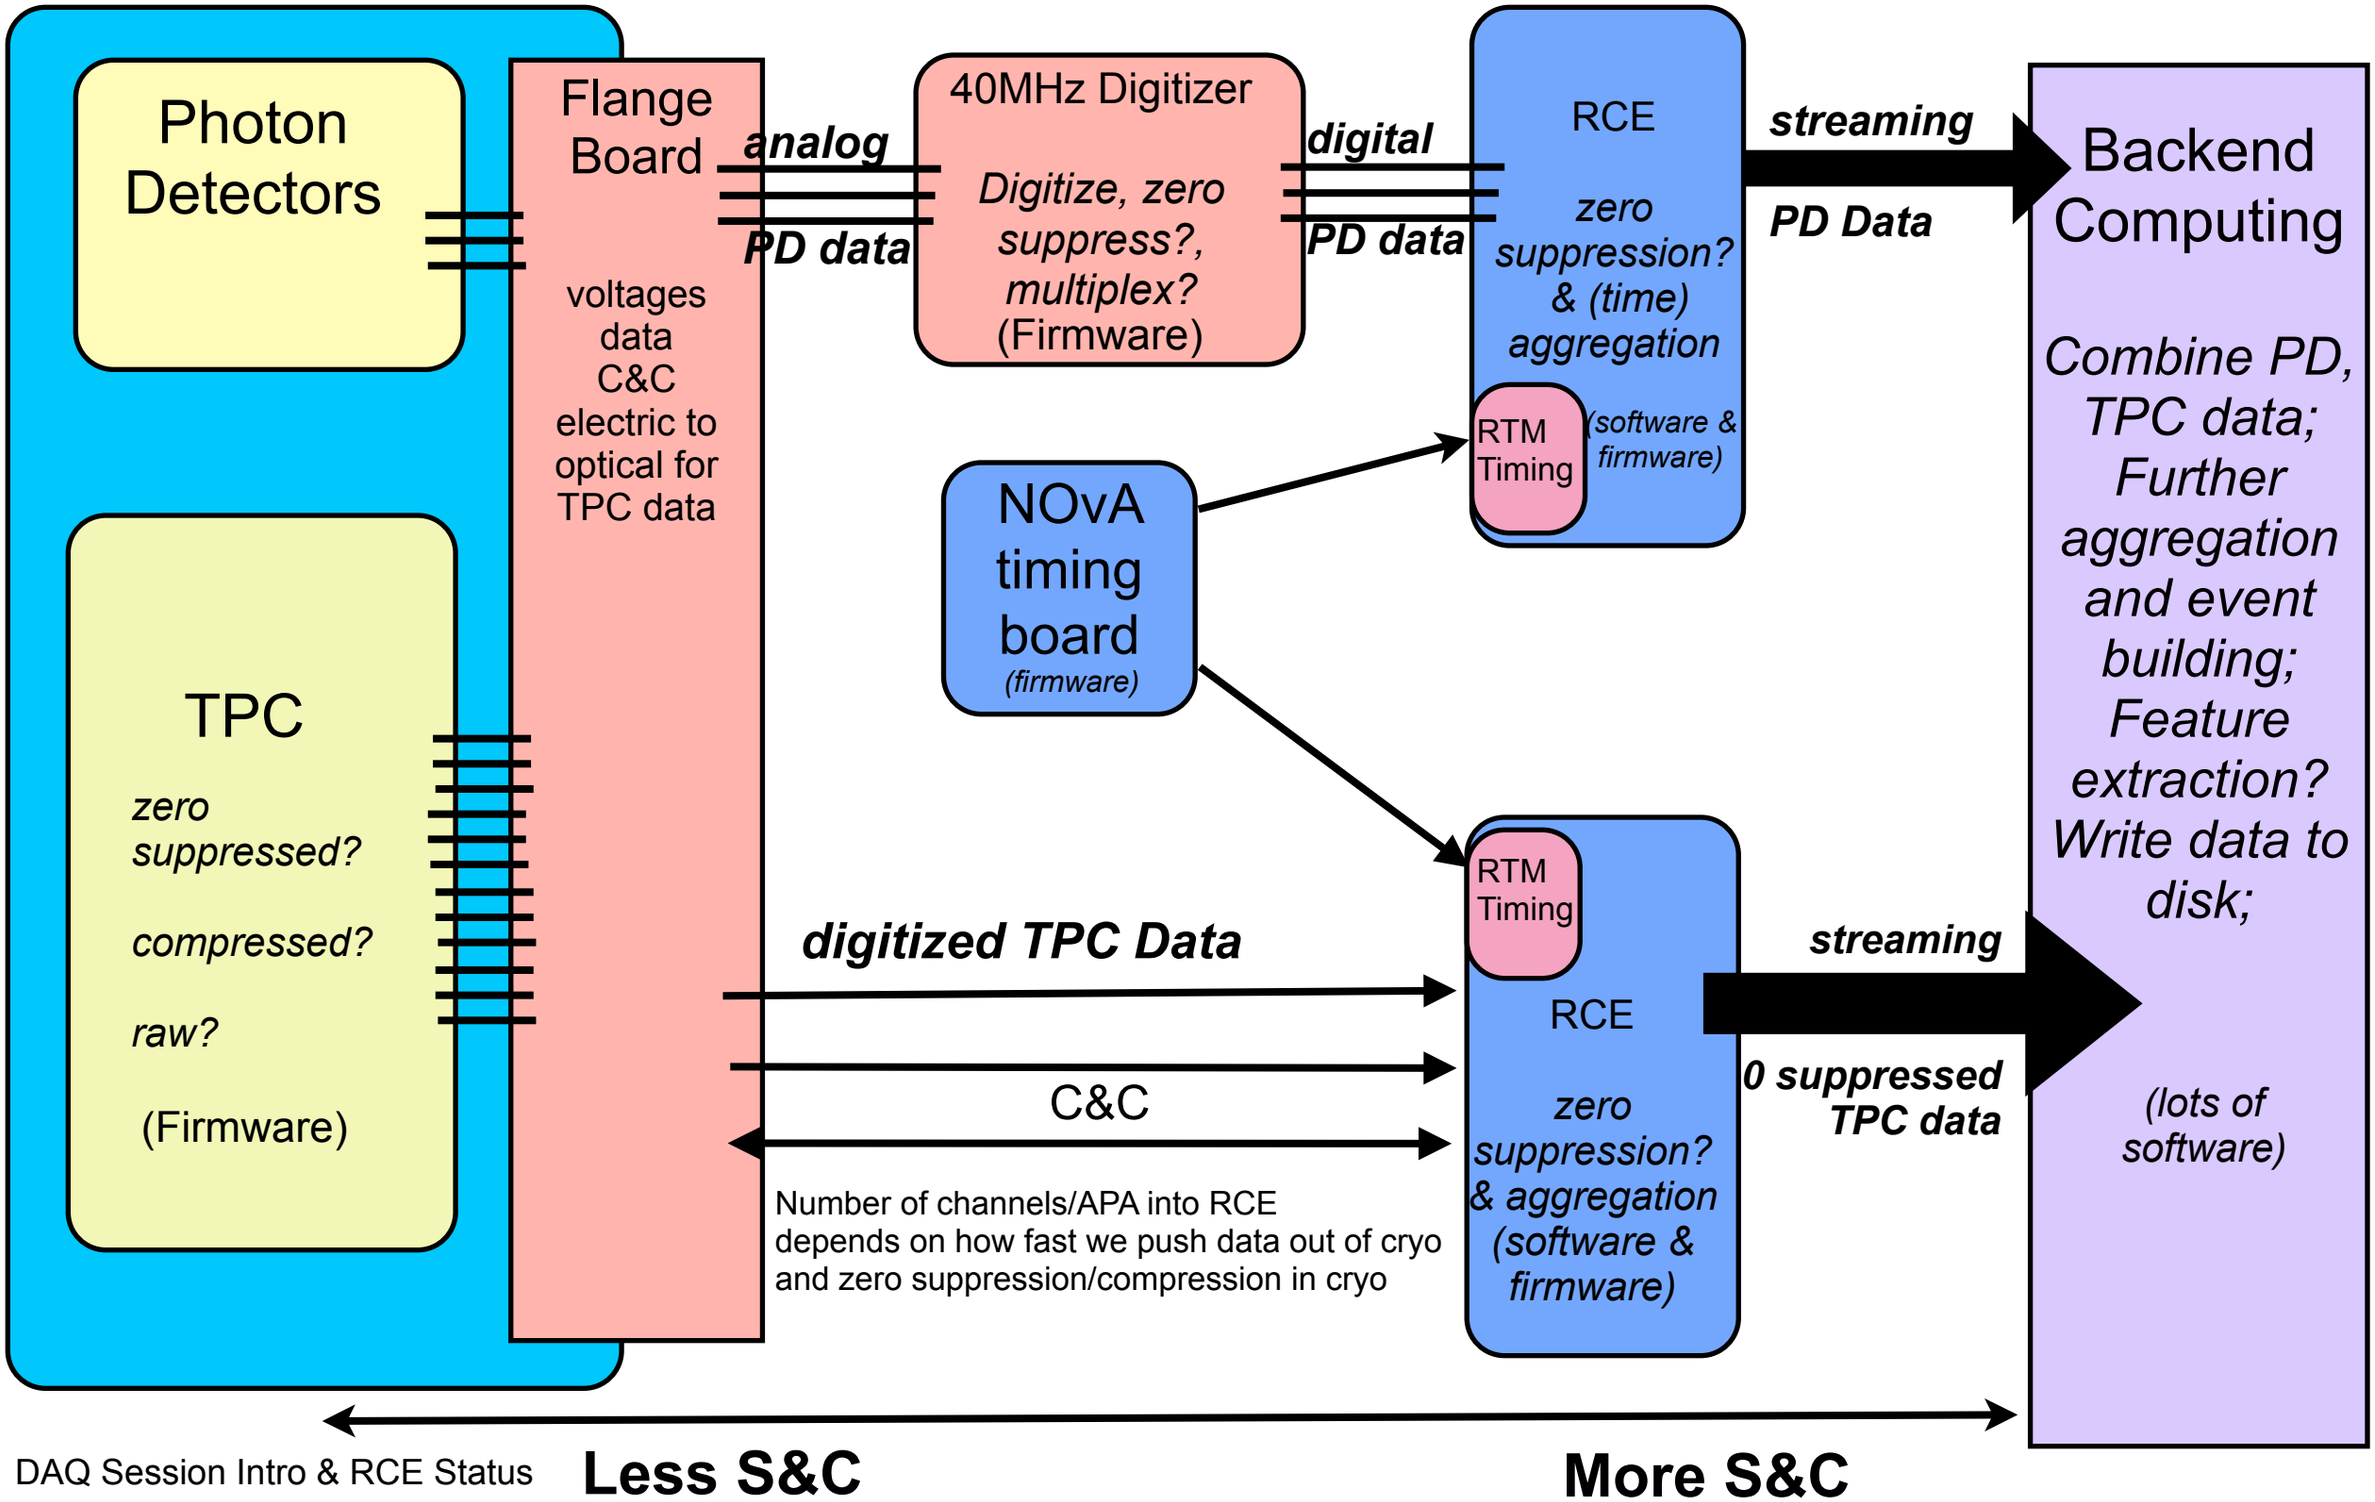

DAQ High-level Block Diagram For 35t

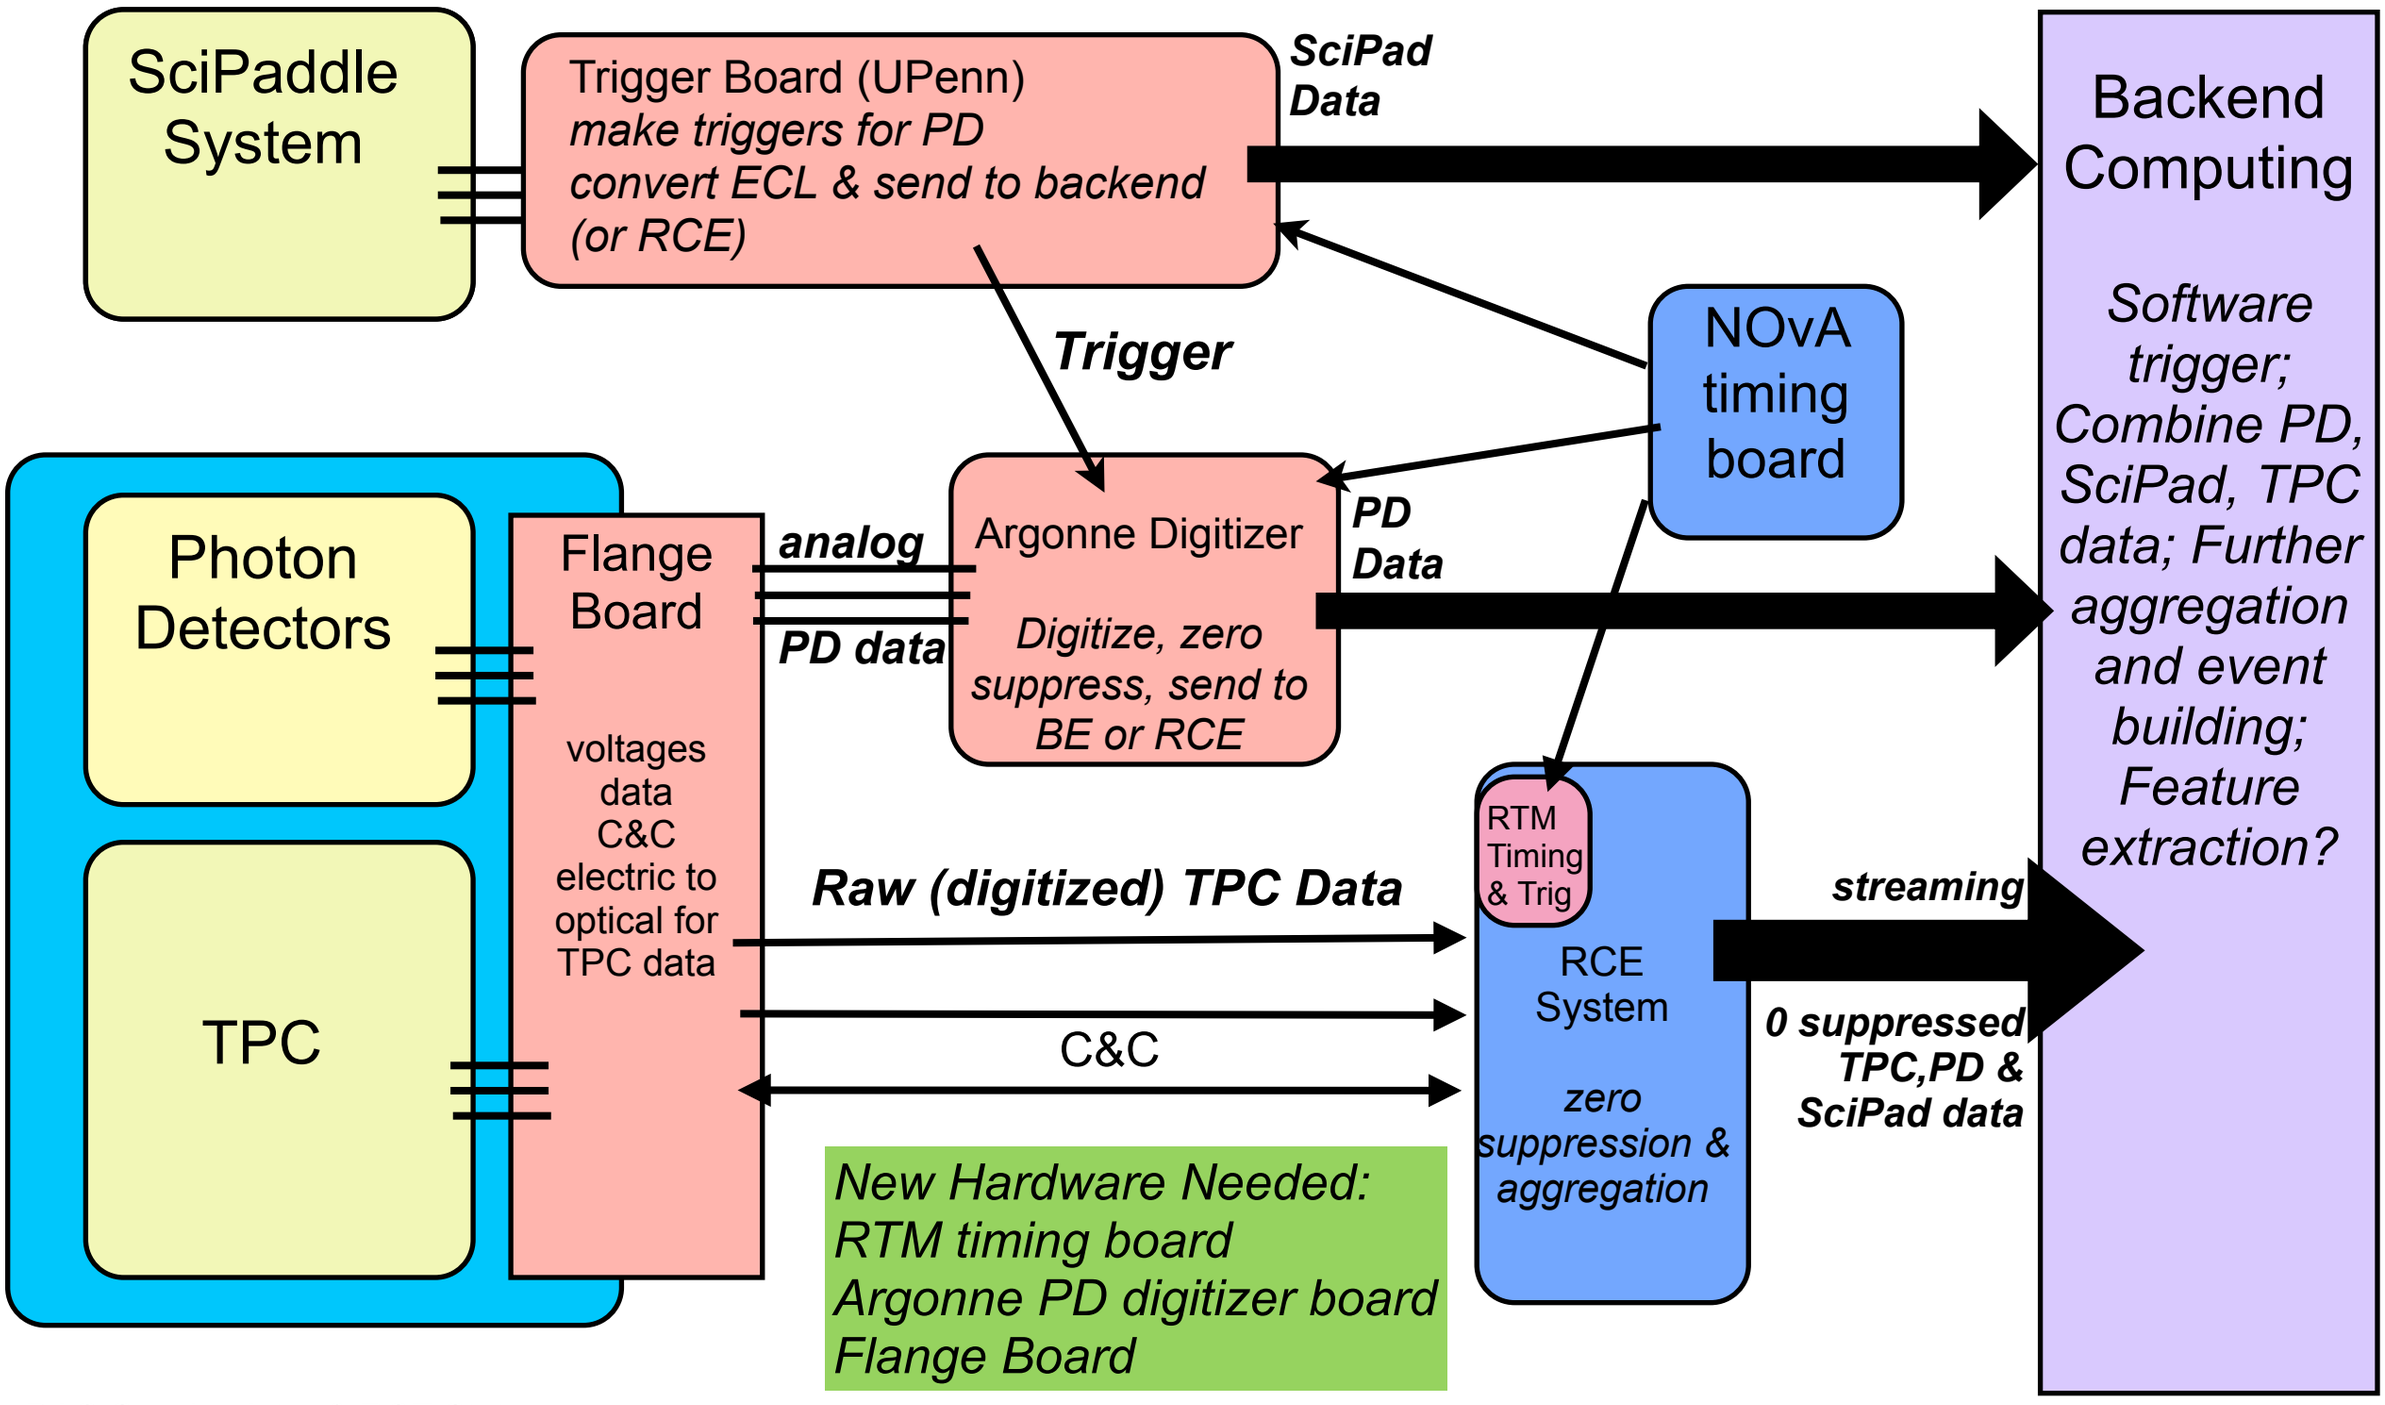

DAQ Workflow

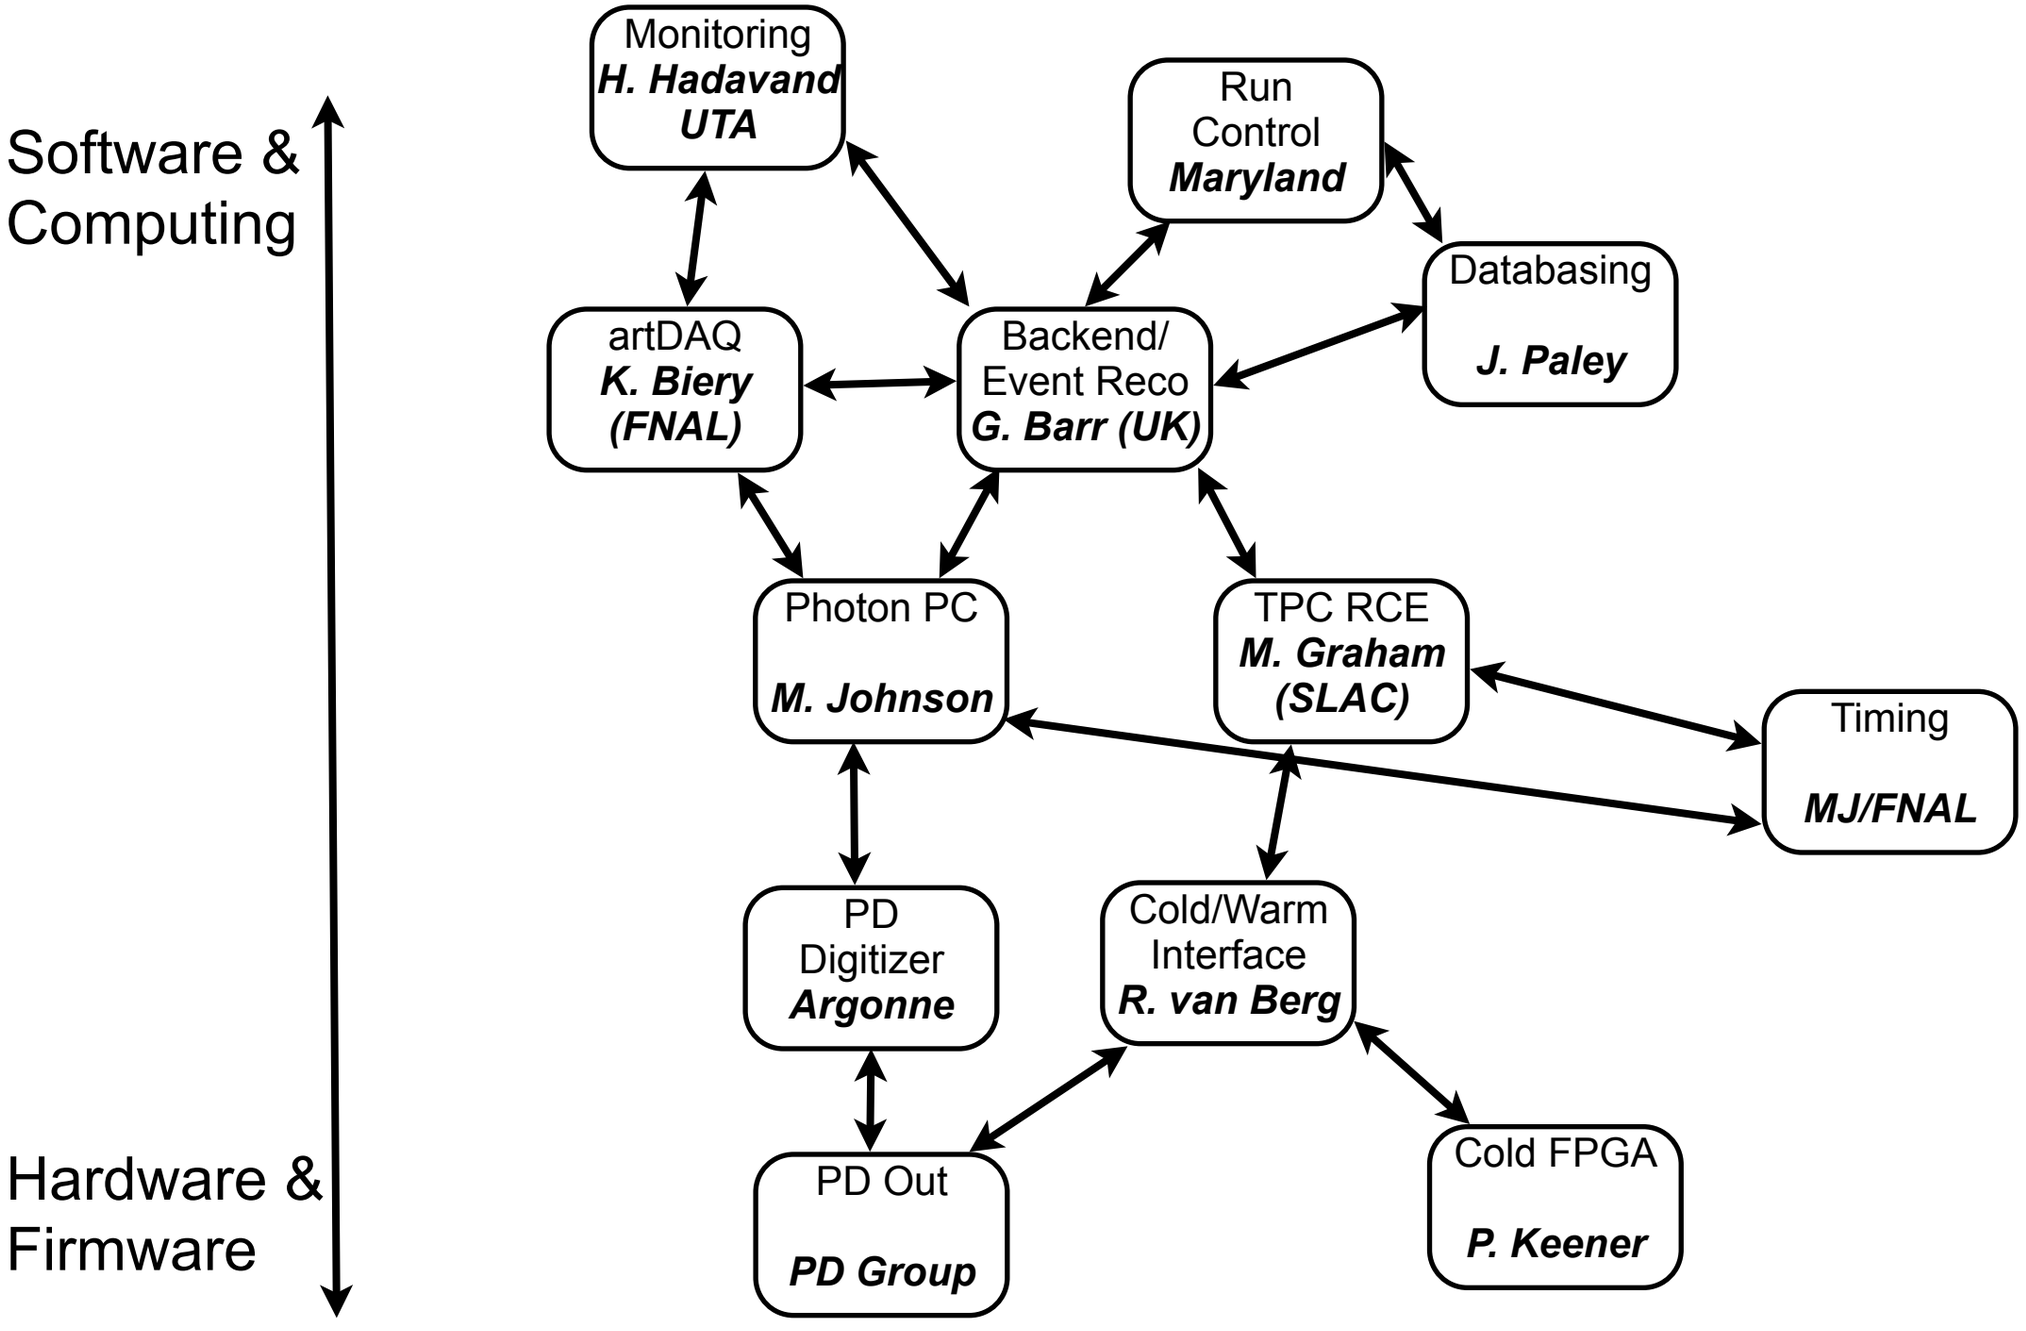

Run Control, Back-end computing, database

- Run Control: Maryland (Erik Blaufuss)
- Back-end computing: G. Barr (Oxford)
  - artDAQ infrastructure: K. Biery (FNAL)
- configuraton database: J. Paley

See talks from CSU parallel sessions:

docdb-7837,7832,7848, 7858

artdaq Components (K. Biery & R. Rechenmacher)

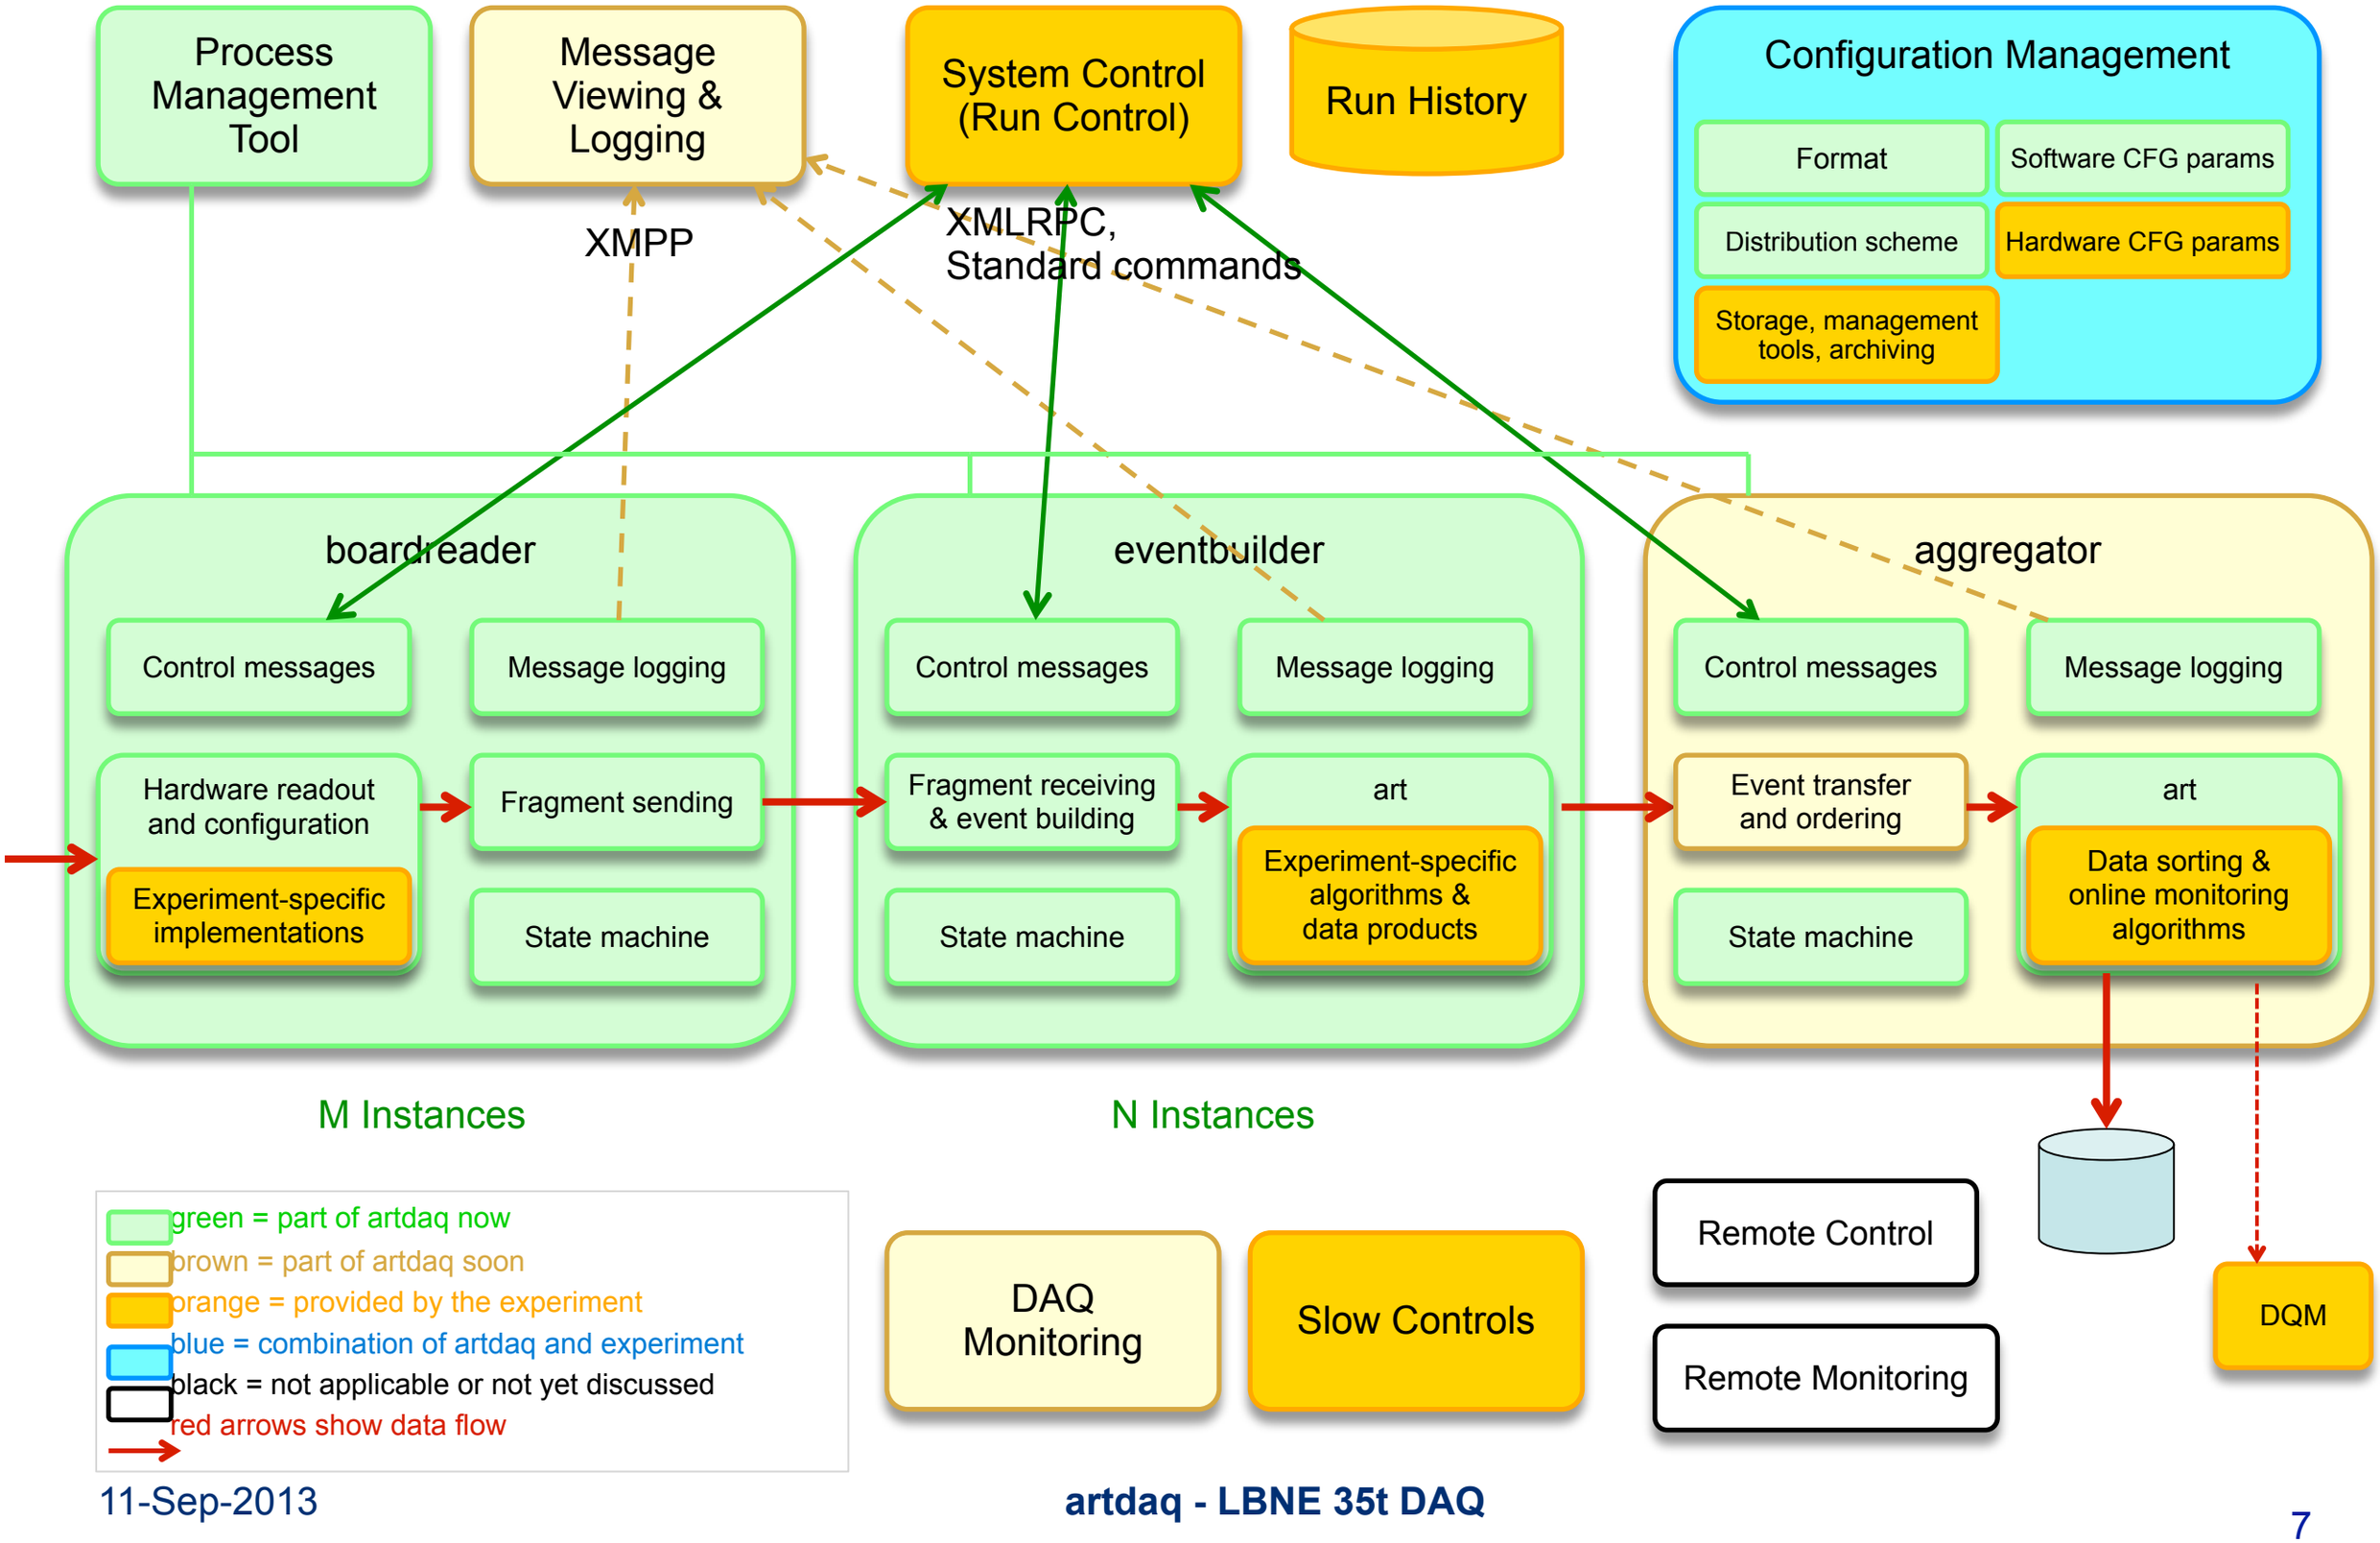

Back-end computing & event building

Giles Barr docdb/7832

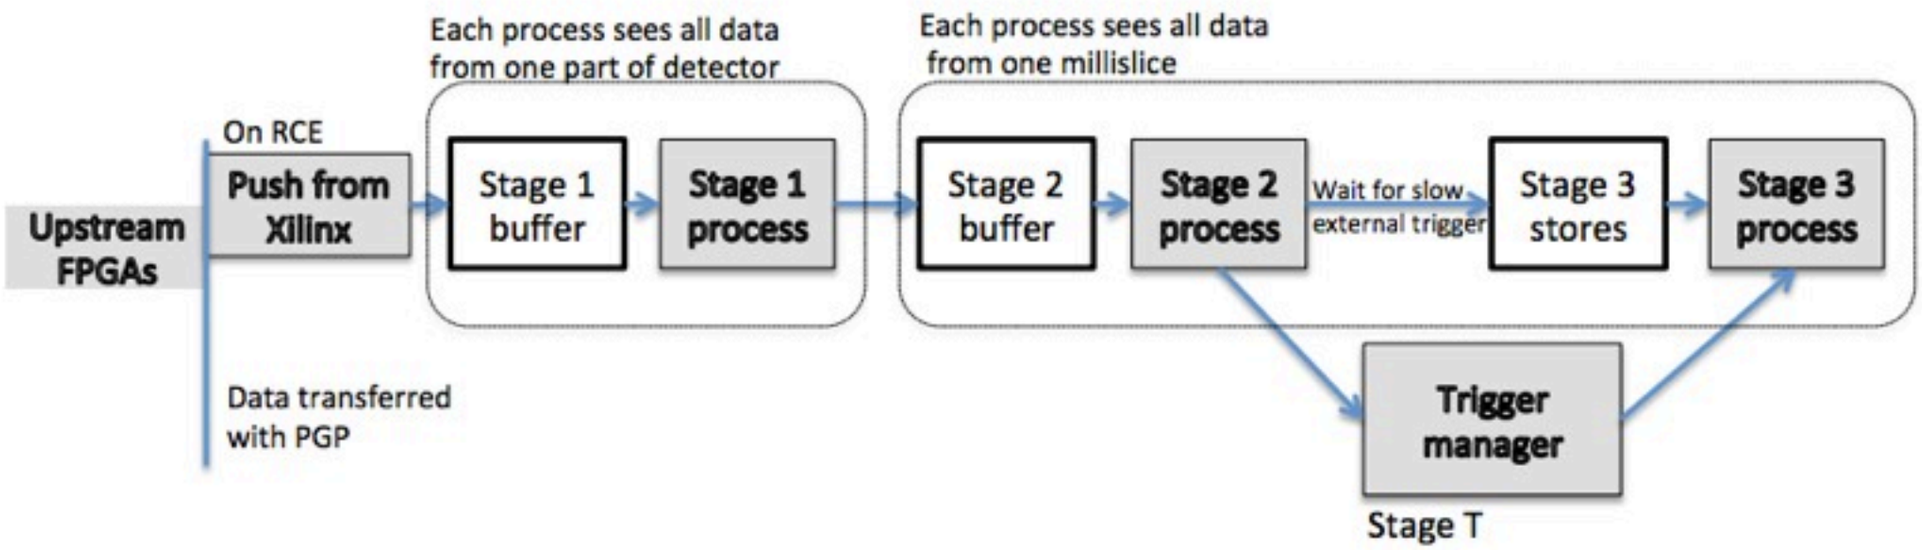

Stage 1	Parallelized by detector section. Receive RCE data, merge micro to milli-slice, route to correct stage 2, per-view trigger processing
Stage 2	Parallelized by milli-slice. With overlaps. SN trigger caching. Complete trigger calculation in milli-slice and overlap.
Stage T	One process in detector, arbitrate spill, calibration, and detector triggers.
Stage 3	Parallelized by milli-slice. Buffer data during trigger decision. Receive trigger decisions, splice milli-slice to event.

Each stage pushes data to the store of the next stage where it waits in queues.

35t rate summary

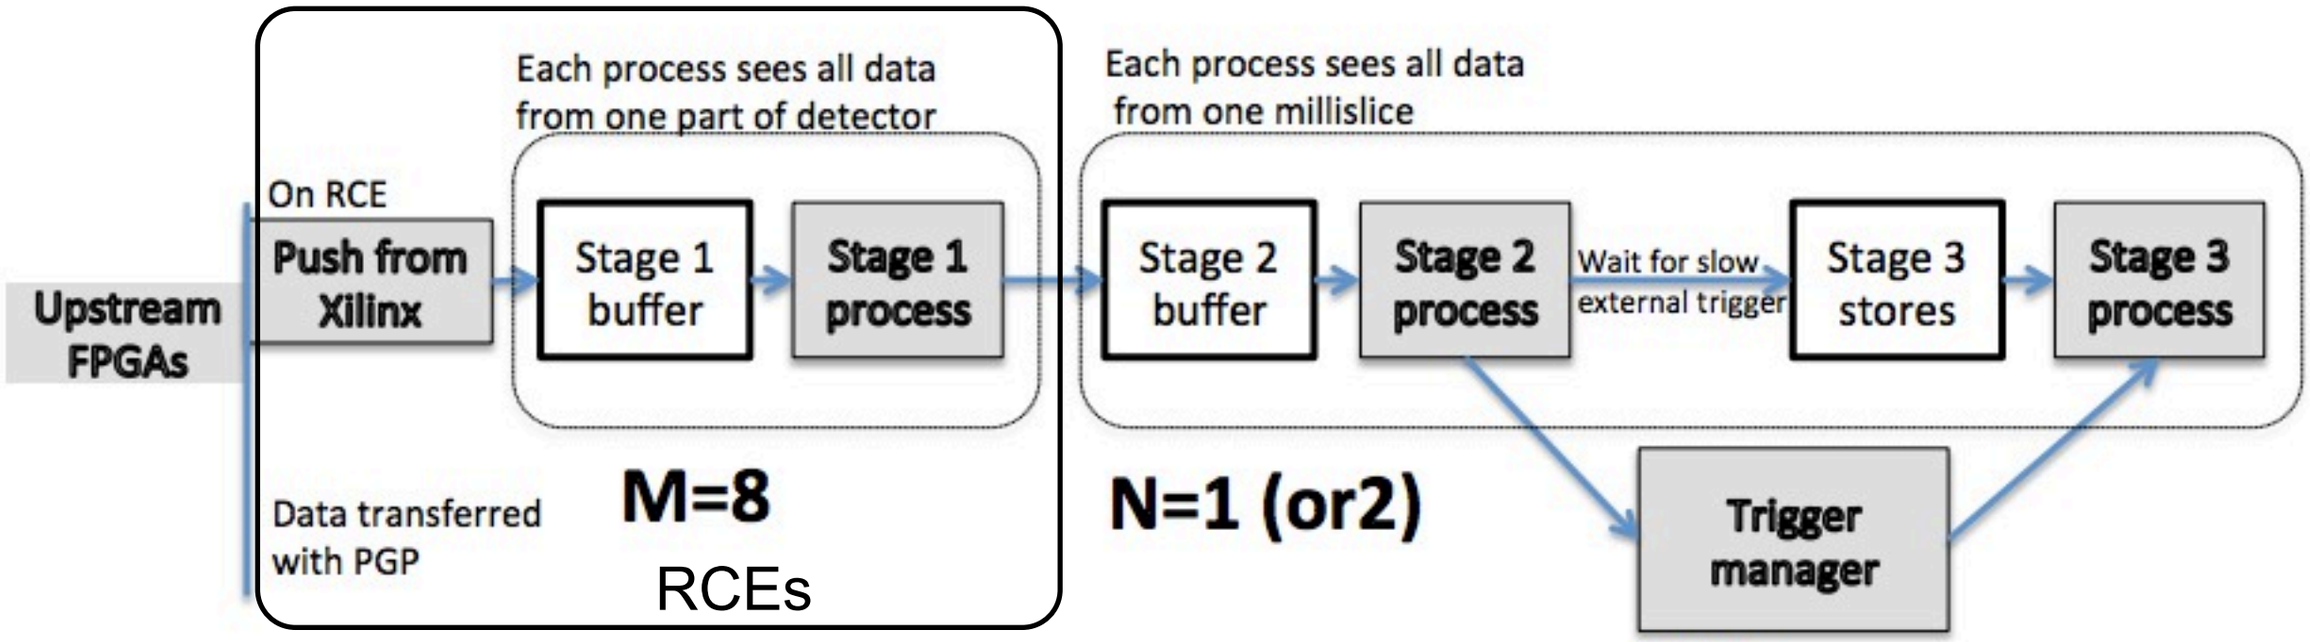

	Rates in Mbytes/s			
	per FEB	per RCE 9	per RCE 8 (12)	full 35t
Non ZS	500	1125	1000 (1500)	9,000
ZS cosmics	2	4.5	4 (6)	34
Radioactivity	0.075	0.2	0.2 (0.3)	1.4
Electronic noise	0.001	0.002	0.002 (0.003)	0.020

(a) 72 fibers, 4 from each CB go to 8 RCEs 9 each. (per RCE 9)

(b) (b) 6 RCEs have 8 fibers (=2 CB), 2 RCEs have 12 fibers (=3 CB).

Can conclude from these rates N=1 stream through stage 2/3 is sufficient.

Change: artDAQ will not be on RCEs, so boardreader will be on COTS computers and will connect to RCE over socket. Once socket is established, RCE still pushes data to board reader.

Giles Barr docdb/7832

With a 40Hz PD&SciPad trigger, this will be <3MB/s to disk

Underground full-LBNE rate summary

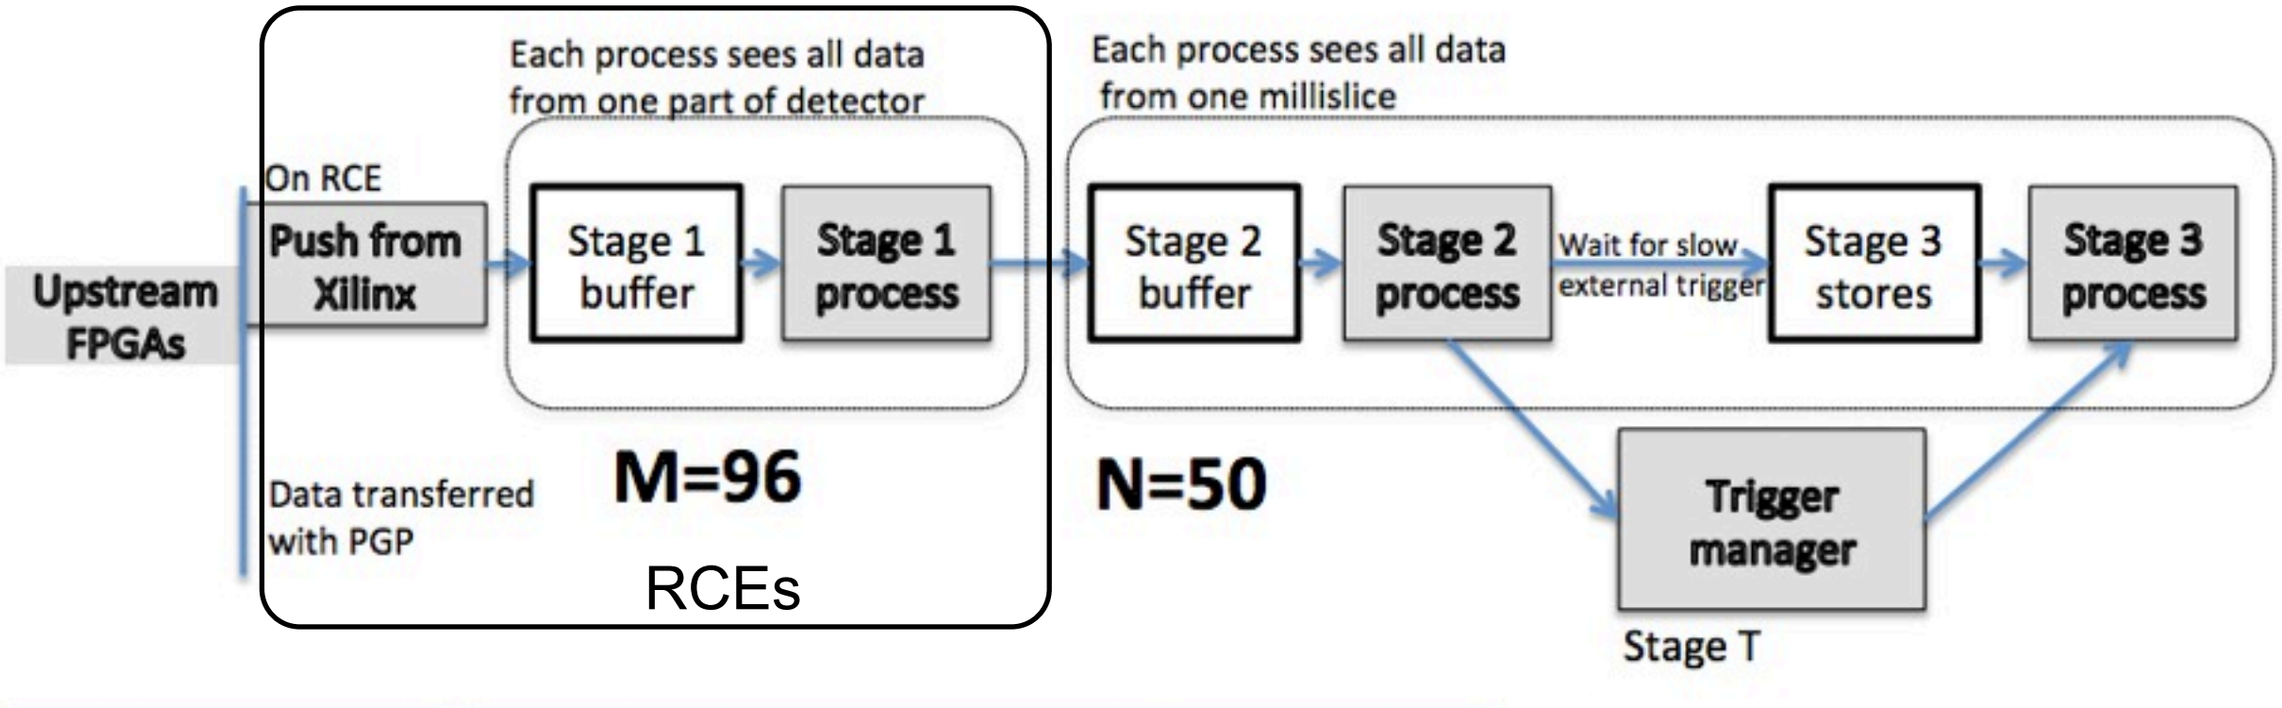

	Rates in Mbytes/s		
	per Stage 1	per Stage 2/3	Full LBNE
Non ZS	(80,000)	-	7,600,000
ZS cosmics	0.05	0.1	5
Radioactivity	10	20	940
Electronic noise	0.2	0.4	20

Rates given here for cosmics are at 800ft level. At 4850ft level, they are even more negligible.

Choose M=96, i.e. one RCE/APA (assumes H/W zero suppr.)

SuperNova store of 6TB disk on each stage2/3 node would give 80 hours of cache.

Keep all cosmics: we get 0.45TB/day

Giles Barr docdb/7832

Surface full-LBNE rate summary

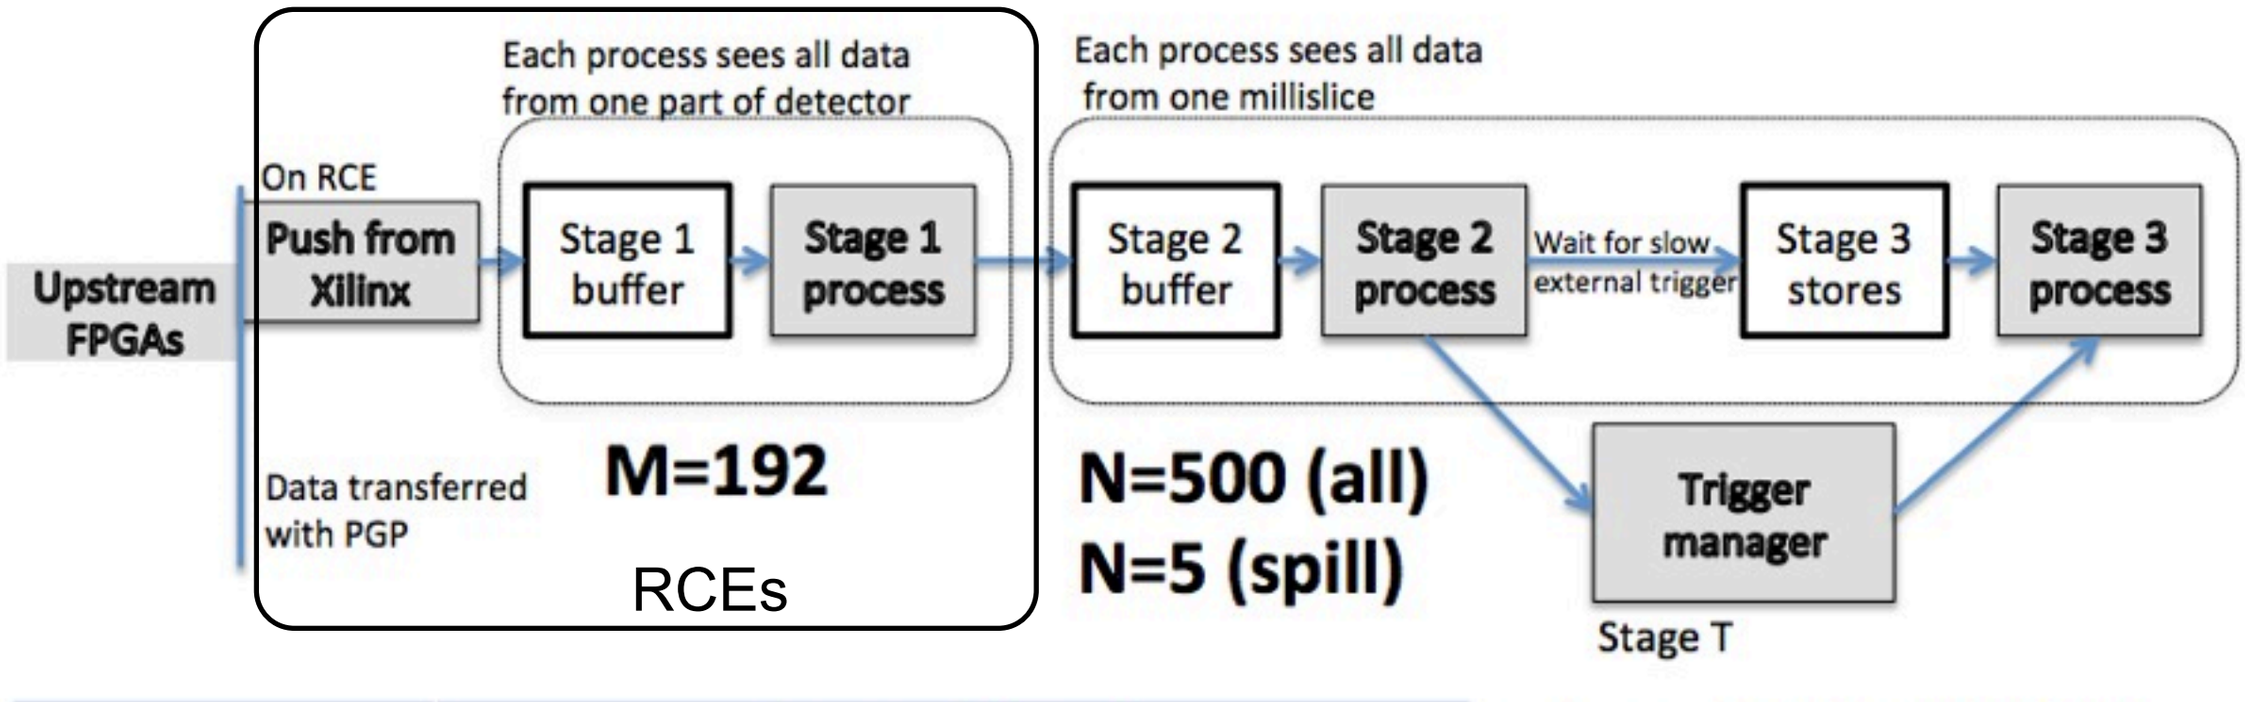

	Rates in Mbytes/s			
	per Stage 1	per Stage 2/3		Full LBNE
		All	Spill	
Non ZS	(40,000)	-	-	7,600,000
ZS cosmics	32	12	6	6,000
Radioactivity	5	2	1	940
Electronic noise	0.1	0.04	0.02	20

Choose M=192 ie 2 RCEs/APA

Spill option is if we take only data around a beam spill 4 x 1.5ms drifts every 1.3sec (factor 200 reduction)

SuperNova store of 6TB disk per each of the 32 stage2/3 nodes (16-cores each) node: Cache of 8.5 hours.

If we write all spills, we get 0.6TB/day.

Giles Barr docdb/7832

Cryostat

FEB

readout & ADC

8 x 16 channel ASICs → 8 out

FPGA

⇒ 3.1 Gbps

~4 Gbps PGP over 4 Cu links

FEB

readout & ADC

8 x 16 channel ASICs → 8 out

FPGA

⇒ 3.1 Gbps

~4 Gbps PGP over 4 Cu links

20 more FEBs

Flange Board

data

Copper ⇒ Fiber

clock/ control

Copper ⇒ Fiber

+ other FEBs

from NOvA timing unit

1:1 Cu:Fiber @ 1Gbps

ATCA Shelf

RTM

system timing daughter board

96 x **1** Gbps

(up to 96 x 3.1Gbps)

COB

RCE

RCE

8 RCEs

up to 14 slots/crate

TCP/IP to Farm

~few 100 Mbps

72 wires/FEB

4: 1Gbps output data
3: clock/command/snych
4: JTAG
5: Low Voltage
3: High Voltage
8: analog
1: test pulse

Front- and Middle-end S&C

- Front-end interactions with software & computing
  - FPGA configuration firmware and FPGA or ASIC parameters: *run-control* initiates commands to load a new config or parameter set from the DAQ database (database may just own links to the config files depending on size...config & parameter IDs will be unique); must propagate down the complete chain
  - no actual software on the front-end
- For the MiddleEnd-->RCE system
  - Configurations served via DAQ database as per FE
  - Additionally, there will be software running on the RCEs doing things like aggregating events, maybe zero suppression, maybe basic pattern recognition...etc
  - RCE-base soft/firmware and LBNE-specific soft/firmware

Summary and Closing Remarks

- Our assignment is to get all of the useful data onto disk, for use of analysts
  - for full LBNE, I think this means we stream (triggerless) the zero suppressed TPC and PD data up to the backend farm (and maybe all to disk as well)
  - for the 35t, I think this is the ideal as well...although it is multi-TB of data/day; we are planning this for the SciPad & TPC...the Argonne digitizer may get us there for the PD as well
- We have a pretty solid plan towards a full DAQ system by Oct. 1 2014...there are places we could use manpower though:
  - LBNE-specific software on the RCEs
  - Software help for backend computing/run control etc would not be turned down
  - Online detector performance monitoring: hit & track recon, event display...using art/LArSoft
    - not sure who's box this falls under
- We're still shaking out just what the backend computing requirements are, even for the 35t...
- my best guess for 35t:
  - 4 cores running artDAQ board readers, evt building, triggering etc (maybe less)
  - dedicated machines running DAQ database and run control
  - a few other machines in the hall for on/offline monitoring
  - XXXTB disk for prompt data storage

Photon Detector & Scintillating Paddle DAQ

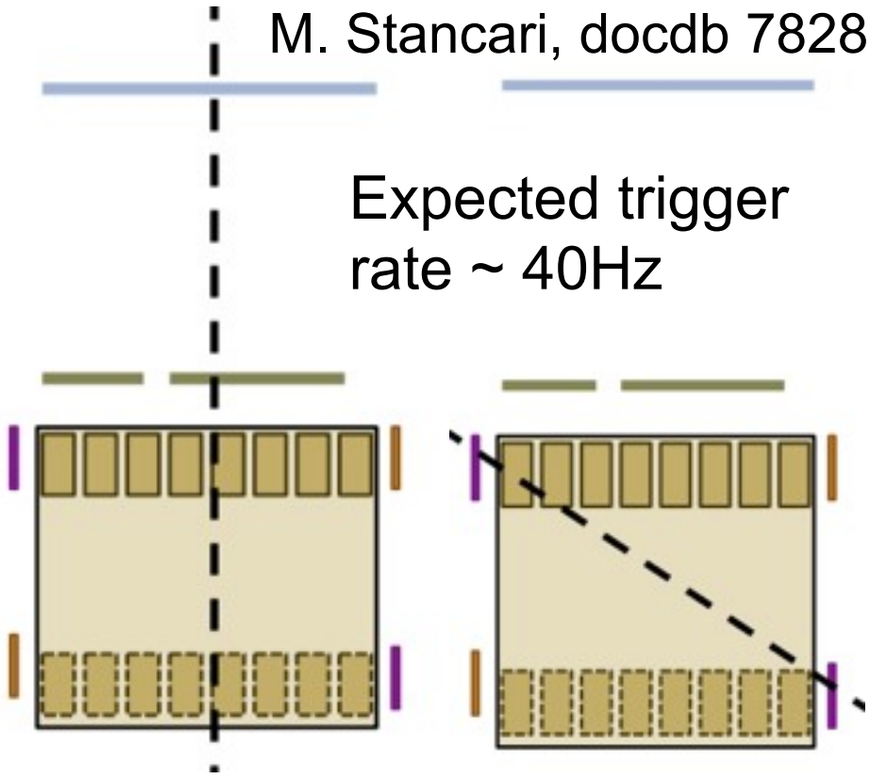

IU "Balloon Boards":

64 channel board "sets"
(we'd use 2 sets)
ADC & TDC ~ 1ns
streaming USB output
(no trigger required)
NIM Trigger out (OR)
(not really what we want... homebrew with NIM units?)

CAEN 5740 digitizer/PC combo

should handle the full 40Hz rate...100Hz likely too much, 10 Hz no problem

NOvA global time is combined at PC via a PCIe card (for mu2e

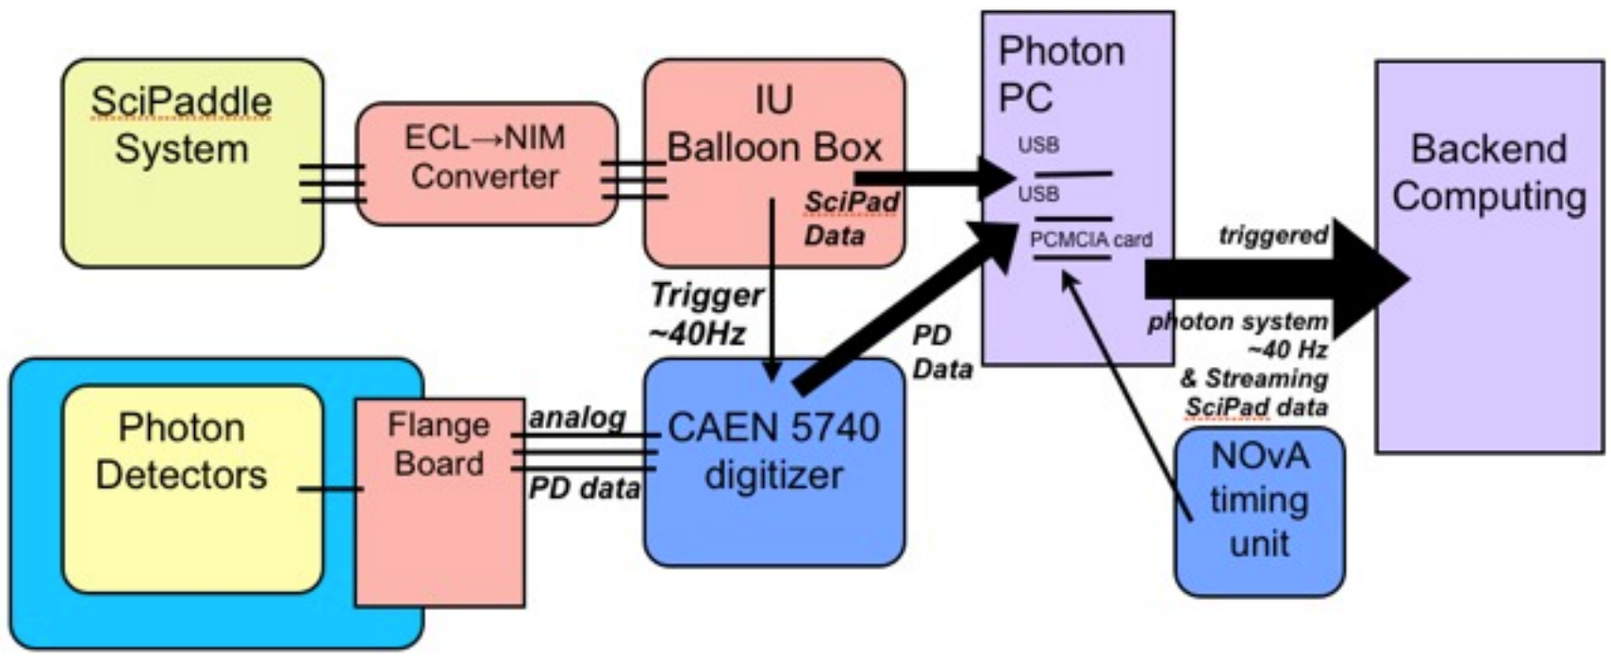

Alternative 35t DAQ High-level Block Diagram (fallback)

SLAC