Data Reduction for protoDUNE/SP TPC

Brett Viren

Physics Department

BROOKHAVEN NATIONAL LABORATORY

S&C

2017 May 17

Outline

Overview of Concept

Reduction Options

Strawman Proposal

Computing Requirements

Processing Latency

Caveat: only considering the **TPC** data here.

Data Reduction Idea in a Nutshell

After performing these required initial steps:

read	Unpack, uncompress and assemble raw *fragments* ($\sim 10\times$ step inflation)
ADC	Correct for ADC nonlinearity and stuck bins.
noise	Remove excess noise, if any.
signal	Apply signal processing: field+electronics response deconvolution, software filter, signal-ROI selection. ($\sim 100\times$ step reduction)

Then, we can save a much smaller output file after applying:

- resampling/rebinning (3 – 4$\times$ step reduction)
- redigitizing/truncating and packing (2$\times$ step, 32bit$\rightarrow$16bit)
- compression (via ROOT I/O) ($\sim 2\times$ step reduction)

Finally, we can send the reduced file on to reconstruction processing.

Details in DocDB 2089

Pictorial

Scheme 1: nominal processing pipeline with **no reduction**.

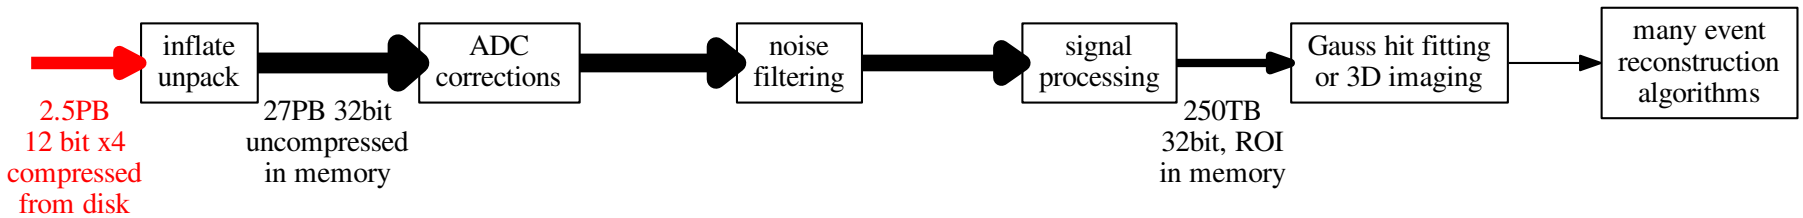

Scheme 2: **including reduction** (for one choice of rebinning).

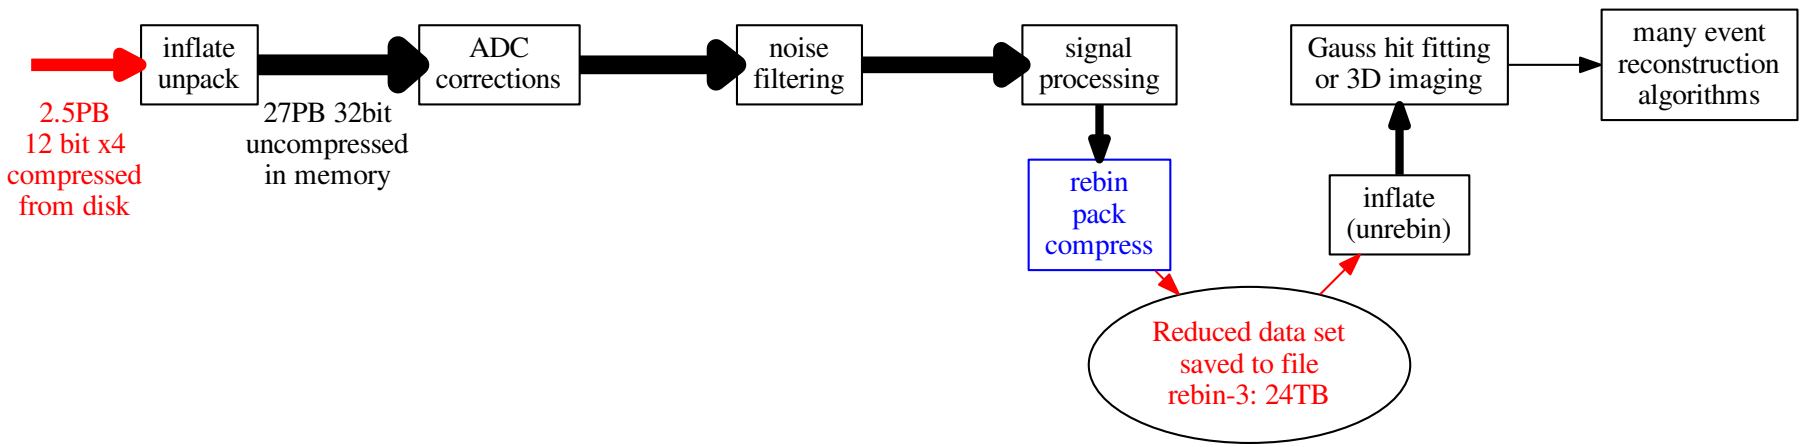

Reduction Efficacy Options

Partial summary of DocDB 2089 study which is based on scaling from MicroBooNE data:

option	total size	note
TPC FEE bit stream	10 PB	triggered
raw, online output:		RCE & FELIX compressed
early μboone noise level	**2.5 PB**	nominal $\times$4 comp
excess μB noise removed	2 PB	possible $\times$6 comp
16bit packed:		ROOT compression
signal-ROI	70 TB	4$\times$ over sampling
rebin-3	24 TB	still over sampled
rebin-4	19 TB	Nyquist limit
32bit packed	$\sim 2\times$	guess, not measured

Proposed Production Run Plan

I propose we do data reduction (scheme #2) in two phases:

1. Begin processing during **commissioning** with **no rebin**.
   - Validate reduction.
   - Early testing of eventual production jobs.
   - Still enjoy some modest reduction.
   - Independently, I think ADC/noise/signal processing algorithms will stabilize during this time period.
2. Transition into **operations** with **rebin-3** or **rebin-4**.
   - Should just entail a job configuration change.

Subsequent **re**processing uses reduced data from 1st processing unless flaws are found in ADC corrections, noise filtering or signal processing. Must then go back to raw.

Computing Requirements

- Data reduction, per se, **requires essentially no CPU** but the previous, required stages (Read, ADC, Noise, Signal) do:
  - 4.5 core·minute/trigger is required.
- Assume 10Hz average pD/SP trigger rate, 60 days running:
  - Average of 25Hz spill trigger and 1-for-1 out-of-spill μ-trigger.

option	requirement	note
full keep up	2700 cores	RAM/node limits may increase this
select keep up	810 cores	if preselect "good" beam triggers
full delayed	4M core·hour	aka: 440 core·year

N.B:

- Above **is just to get through reduction**. Reconstruction takes 30-45 core·minute/trigger or $5-10\times$ more.
- "Select keep up" means relying on Beam Information to identify triggers of likely desired PID **prior** to running any processing and then **delaying processing** of all other beam triggers and all μ-triggers. Estimated 30% reduction (Junk/Norman).

How fast can we process data?

Full chain pD/SP needs (4.5 minute SP + $>$ 30 minute reco)/trigger

- Now'ish: CERN Tier-0: 1500 cores, FNAL Tier-1: 1000 cores (R.Sulej)
  - Shared for many uses in DUNE, both SP & DP.
- Statement in *DUNE Computing Plan In Support of the ProtoDUNE Program Executive Summary*: "Routine '**keep-up**' processing, which involves the first stages of data processing and reconstruction, will be run on FNAL compute resources".
  - Does "keep-up" here mean fixed latency between trigger and result?
  - Does it mean "full" or "select"?
- 2700 cores to get up through data reduction in "keep up" mode.
  - Here, "keep up" indeed means fixed trigger–result latency.
  - We need another **18,000 cores for reconstruction** to run in "keep up" mode.
  - If use one year for 1st-pass full chain pD/SP processing: 3400 cores.

How fast can we process data?

Full chain pD/SP needs (4.5 minute SP + $>$ 30 minute reco)/trigger

- Now'ish: CERN Tier-0: 1500 cores, FNAL Tier-1: 1000 cores (R.Sulej)
  - Shared for many uses in DUNE, both SP & DP.
- Statement in *DUNE Computing Plan In Support of the ProtoDUNE Program Executive Summary*: "Routine '**keep-up**' processing, which involves the first stages of data processing and reconstruction, will be run on FNAL compute resources".
  - Does "keep-up" here mean fixed latency between trigger and result?
  - Does it mean "full" or "select"?
- 2700 cores to get up through data reduction in "keep up" mode.
  - Here, "keep up" indeed means fixed trigger–result latency.
  - We need another **18,000 cores for reconstruction** to run in "keep up" mode.
  - If use one year for 1st-pass full chain pD/SP processing: 3400 cores.

How fast do we need to process?
How many reprocessing per year?
How many cores will pD/SP have?

How fast can we process data?

Full chain pD/SP needs (4.5 minute SP + $>$ 30 minute reco)/trigger

- Now'ish: CERN Tier-0: 1500 cores, FNAL Tier-1: 1000 cores (R.Sulej)
  - Shared for many uses in DUNE, both SP & DP.
- Statement in *DUNE Computing Plan In Support of the ProtoDUNE Program Executive Summary*: "Routine '**keep-up**' processing, which involves the first stages of data processing and reconstruction, will be run on FNAL compute resources".
  - Does "keep-up" here mean fixed latency between trigger and result?
  - Does it mean "full" or "select"?
- 2700 cores to get up through data reduction in "keep up" mode.
  - Here, "keep up" indeed means fixed trigger–result latency.
  - We need another **18,000 cores for reconstruction** to run in "keep up" mode.
  - If use one year for 1st-pass full chain pD/SP processing: 3400 cores.

How fast do we need to process?
How many reprocessing per year?
How many cores will pD/SP have?

I don't know