Data Reduction Proposal and Computing Requirements

Brett Viren

Physics Department

BROOKHAVEN
NATIONAL LABORATORY

pD-SP DAQ / S&C
DUNE Collaboration Meeting
25 Jan 2017

Outline

Proposal for Initial Data Reduction Process

Data Formatting Issues

Additional Opportunities

Reduction Proposal

- Run a **data reduction process** soon after raw data is saved.
  - Algorithms (noise & signal processing) must run somewhere in any case.
  - Strategy: "do no harm to information content".
- Save reduced data and **still save full raw data to tape!**
  - Eg, for understanding detector during commissioning.
  - Any special analysis can also still run on full raw data.
- Later, improve reduction? $\Rightarrow$ re-run on full raw data.

Some uses for the **highly reduced** data set:

- **Efficiently feed production** jobs ("keep up" and re-re-re-processing).
- **Easy sharing of fresh data** to look for/at **emerging detector issues**.
- Tidy samples **facilitate reco development**.
- **Everyone can have full copy** of all signal waveforms on their workstation!
- Maybe facilitate some **public outreach / data sharing**?

Many technical details are in DocDB 2089

Reduction Methods, Options and Evaluation

evaluation	use a few uBoone raw events (`ADC` tier)
scaling	protoDUNE has $\sim 2\times$ channels of uBoone
noise filter	reduces entropy, helps compression (`SUB` tier)
	$\rightarrow$ if all goes well, emulates expected protoDUNE data
ROI selection	keep only **signal regions of interest** (`ROI` tier)
resample	sum N neighboring bins to **remove oversampling** (`RBn` tiers)
	$\rightarrow$ but do not violate Nyquist theorem!
truncation	32 bit float working samples to 16 bit int (for `ROI` or `RBn`)
packing	12 (for `ADC`/`SUB`), 16 or 32 bit packing into 32 bit output arrays.
compression	apply ROOT compression (Z level of 0-9)

Look at **reduction factor** and **compression time** as a function of all meaningful combinations of above options $\longrightarrow$

Reduction Options and Efficacy

tier	bits	Z-lvl	Z-time	factor	2.5 PB
ADC	12	5	6.8 sec	**4.2**	0.6 PB
ADC	12	9	**4.3 min**	5.5	0.5 PB
SUB	12	5	5.6 sec	6.0	0.4 PB
SUB	12	9	**3.6 min**	8.0	0.3 PB
ROI	16	9	0.5 sec	150	17 TB
RB2	16	9	0.3 sec	260	9.6 TB
RB3	16	9	0.2 sec	415	**6.0 TB**
RB4	16	9	0.1 sec	540	4.6 TB

(just a summary, see note for more)

- Noise filtering and signal processing takes **4.5 min/trig** (6 APA)
  - → Z-time above is scaled to 6 APA protoDUNE channel count.
- 12 bits means fully packing 8 samples across 3 words (32 bit int).
- Can **achieve expected factor of 4 raw data compression** at uBoone noise level.
  - → expect factor 6-8 raw data compression if low-noise, but Z-time grows fast!
- Z-time not significant once some reduction is included.
- This **only counts the TPC data**, other data sizes may become significant!
- RB4 rebins to exactly the Nyquist limit, **RB3 is the safe recomendation**.

Data Formatting Issues

Expect that:

- Reduction process directly reads raw data files into working memory.
- Any fragment types not touched are saved to output as-is.
  - (just nominal plan, see next slide for opportunities)
- Write reduced data, packed similarly to raw data (ie, `artdaq::Fragment`)

$\Rightarrow$ develop a "reduced DAL" which is sister to the "raw DAL"[1]

[1] Data Access Library, some more details here

Additional Opportunities?

- Basic proposal addresses only TPC data.
- Process provides opportunity for additional processing.
- Best to not require substantially more CPU/RAM.
- Maybe:
  - Merge in any out-of-band Beam Info?
  - Data reduction on non-TPC fragment types?
  - Any reformatting or other "convenience" operations that help downstream read non-TPC data?

If there are proponents for any ideas, let's discuss!

CPU and RAM Estimation

- LArSoft RAM overhead assumed 1–2 GB. Reduction adds negligible if 1 APA at a time, o.w. worse case: +1.2 GB.
- CPU time: 45 sec/APA/trig $\equiv$ 4.5 min/pD-SP/trig.
- 10 Hz trig $\Rightarrow$ 2700 cores to keep up with beam cycle
  - or 6750 cores to keep up with in-spill 25 Hz
- Assume computer = 2 CPUs / 20 cores $\rightarrow$ 135 computers with 64 GB RAM each! (at worse case)
- $\therefore$ Mitigate RAM pressure: use plain-art / stand-alone code, hold only 1 APA in memory, increase node count to 150:

To keep up with beam cycle we need to allocate:

3000 cores with 2 GB RAM each

- **Exclusive CPU** but only during actual protoDUNE/SP running.
- Need to better understand the **RAM requirement**.
- Any additional **processing beyond TPC** needs understanding.